arXiv:1604.04707v2 [cs.CG] 22 Apr 2016

Approximation Algorithms for Max-Morse Matching

Abhishek Rathod*[1], Talha Bin Masood†[1] and Vijay Natarajan‡[1]

[1]Department of Computer Science and Automation, Indian Institute of Science, Bangalore, India

April 25, 2016

Abstract

In this paper, we prove that the Max-Morse Matching Problem is approximable, thus resolving an open problem posed by Joswig and Pfetsch [15]. For D-dimensional simplicial complexes, we obtain a $(D+1)/(D^2+D+1)$-factor approximation ratio using a simple edge reorientation algorithm that removes cycles. We also describe a $2/D$-factor approximation algorithm for simplicial manifolds by processing the simplices in increasing order of dimension. This algorithm may also be applied to non-manifolds resulting in a $1/(D+1)$-factor approximation ratio. One application of these algorithms is towards efficient homology computation of simplicial complexes. Experiments using a prototype implementation on several datasets indicate that the algorithm computes near optimal results.

1 Introduction

Discrete Morse theory is a combinatorial analogue of Morse theory that is applicable to cell complexes [9]. It has become a popular tool in computational topology and visualization communities [6, 26] and is actively studied in algebraic, geometric, and topological combinatorics [17, 23].

The idea of using discrete Morse theory to speedup homology [12], persistent homology [24] and multidimensional persistence [1] computations hinges on the fact that discrete Morse theory helps reduce the problem of computing homology groups on an input simplicial complex to computing homology groups on a collapsed cell complex. Ideally, if one were to compute a discrete gradient vector field with minimum number of critical simplices (unmatched vertices in the Hasse graph) or maximum number of regular simplices (matched Hasse graph vertices), then the time required for computing homology over the collapsed cell complex would be the smallest. However, finding a vector field with maximum number of gradient pairs is an NP-hard problem as observed by Lewiner [18] and Joswig et.al. [15] by showing a reduction from *the collapsibility problem* introduced by Eğecioğlu and Gonzalez in [8]. We study the problem of efficiently computing an approximation to the maximum number of gradient pairs in a discrete gradient vector field.

Computing the homology groups has several applications, particularly, in material sciences, imaging, pattern classification and computer assisted proofs in dynamics [16]. More recently, homology and persistent homology have been appraised to be a more widely applicable computational invariant of topological spaces, arising from practical data sets of interest [5]. An approximately optimal Morse matching computed using the algorithms described in this paper may be used towards efficient computation of homology. One of the primary motivations for us to initiate the study of approximation algorithms for discrete Morse theory was that a previous study [12] involving discrete Morse theory in homology computation reported noteworthy speedup over existing methods. Their method used a modification of the coreduction heuristic [25] to construct discrete Morse functions. We start with a twin goal in mind – first to introduce rigour into the study by developing algorithms with approximation bounds and then to have a practical implementation that achieves nearly optimal solutions.

*abhishek@jcrathod.in

†tbmasood@csa.iisc.ernet.in

‡vijayn@csa.iisc.ernet.in

1.1 Max Morse Matching Problem

The Max Morse Matching Problem (MMMP) can be described as follows: Consider the Hasse graph $\mathcal{H}_\mathcal{K}$ of a simplicial complex $\mathcal{K}$ whose edges are all directed from a simplex to its lower dimensional facets. Associate a matching induced reorientation to $\mathcal{H}_\mathcal{K}$ such that the resulting oriented graph $\overline{\mathcal{H}_\mathcal{K}}$ is acyclic. The goal is to maximize the cardinality of matched (regular) nodes. Equivalently, the goal is to maximize the number of gradient pairs. The approximate version of Max Morse Matching Problem seeks an algorithm that computes a Morse Matching whose cardinality is within a factor α of the optimal solution for every instance of the problem.

1.2 Prior work

Joswig et al. [15] established the NP-completeness of Morse Matching Problem. They also posed the approximability of Max Morse Matching as an open problem pg. 6 Sec. 4 [15]. Several followup efforts seek optimality of Morse matchings either by restricting the problem to 2-manifolds or by applying heuristics [2, 3, 12, 13, 15, 19, 20]. Recently, Burton et al. [4] developed an FPT algorithm for designing optimal Morse functions.

1.3 Summary of results

We describe a $^{(D+1)}/_{(D^2+D+1)}$-factor approximation algorithm for Max Morse Matching Problem on D-dimensional simplicial complexes. This algorithm uses maximum cardinality bipartite matching on the Hasse graph $\mathcal{H}_\mathcal{K}$ to orient it. We then use a BFS-like traversal of the oriented Hasse graph $\overline{\mathcal{H}_\mathcal{K}}$ to classify matching edges as either forward edges if they do not introduce cycles or backward edges if they do. We then use a counting argument to prove an approximation bound that holds for manifold as well as non-manifold complexes.

For simplicial manifolds, we propose two approximation algorithms that exploit the multipartite structure of the Hasse graph. The first approximation algorithm provides a ratio of $^2/_{(D+1)}$. The ratio is improved to $^2/_D$ via a refinement that specifies the order in which the graph is processed. Both algorithms process simplices of lowest dimension first and then move onto increasingly higher dimensions. Every d-dimensional simplex is first given the opportunity to match to a $(d-1)$-dimensional simplex. If unsuccessful, it is then given the option of matching to a $(d+1)$-dimensional simplex. Furthermore, both algorithms employ optimal algorithms for designing gradient fields for 0-dimensional and D-dimensional simplices (in case of manifolds). The refinement processes subgraphs with small vertex degree with higher priority and hence achieves the better approximation ratio.

We provide evidence of practical utility of our algorithms through an extensive series of computational experiments.

2 Background

2.1 Discrete Morse theory

Our focus in this paper is limited to simplicial complexes and hence we restrict the discussion of Forman's Morse theory below to simplicial complexes. Please refer to [10] for a compelling expository introduction.

Let $\mathcal{K}$ be a simplicial complex and let σ^d, τ^{d-1} be simplices[1] of $\mathcal{K}$. The relation $\prec$ is defined as: $\tau \prec \sigma \Leftrightarrow \{\tau \subset \sigma \text{ and } \dim\tau = \dim\sigma - 1\}$. Alternatively, we say that τ is the *facet* of σ and σ is a *cofacet* of τ. The boundary $bd(\sigma)$ and the coboundary $cbd(\sigma)$ of a simplex are defined as: $bd(\sigma) = \{\tau | \tau \prec \sigma\}$ and $cbd(\sigma) = \{\rho | \sigma \prec \rho\}$. A function $f : \mathcal{K} \to \mathbb{R}$ is called a *discrete Morse function* if it assigns higher values to cofacets, with at most one exception at each simplex. Specifically, a function $f : \mathcal{K} \to \mathbb{R}$ is a discrete Morse function if for every $\sigma \in \mathcal{K}$, $\mathcal{N}_1(\sigma) = \#\{\rho \in cbd(\sigma) | f(\rho) \leq f(\sigma)\} \leq 1$ and $\mathcal{N}_2(\sigma) = \#\{\tau \in bd(\sigma) | f(\tau) \geq f(\sigma)\} \leq 1$. If $\mathcal{N}_1(\sigma) = \mathcal{N}_2(\sigma) = 0$ then the simplex σ is *critical*, else it is *regular*.

A pair of simplices $\langle \alpha^m, \beta^{(m+1)} \rangle$ with $\alpha \prec \beta$ and $f(\alpha) \geq f(\beta)$ determines a *gradient pair*. Each simplex must occur in at most one gradient pair of $\mathcal{V}$. A *discrete gradient vector field* $\mathcal{V}$ corresponding to a discrete

[1] An d-dimensional simplex σ^d may be denoted either as σ or σ^d depending on whether we wish to emphasize its dimension.

Morse function f is a collection of simplicial pairs $\left\langle \alpha^{(p)}, \beta^{(p+1)} \right\rangle$ such that $\left\langle \alpha^{(p)}, \beta^{(p+1)} \right\rangle \in \mathcal{V}$ if and only if $f(\alpha) \geq f(\beta)$.

A simplicial sequence $\{\sigma_0^m, \tau_0^{m+1}, \sigma_1^m, \tau_1^{m+1}, \ldots \sigma_q^m, \tau_q^{m+1}, \sigma_{q+1}^m\}$ consisting of distinct simplices $(\sigma_i \prec \tau_i) \in \mathcal{V}$ and $\sigma_{i+1} \prec \tau_i$ is called a *gradient path* of f.

2.2 The Hasse graph of a simplicial complex

The *Hasse graph* $\mathcal{H}_\mathcal{K}$ of a simplicial complex $\mathcal{K}$ is an undirected graph whose vertices are in one-to-one correspondence with the simplices of the complex. To every simplex $\sigma_\mathcal{K}^d \in \mathcal{K}$ associate a vertex $\sigma_\mathcal{H}^d \in \mathcal{H}_\mathcal{K}$. The edges in the Hasse graph are determined by facet incidences. $\mathcal{H}_\mathcal{K}$ contains an edge between a vertex that represents[2] simplex σ^d and a vertex that represent simplex τ^{d-1} if and only if $\tau \prec \sigma$.

We refer to the set of vertices in $\mathcal{H}_\mathcal{K}$ representing d-dimensional simplices as the *d-level* of the Hasse graph. The *d-interface* of $\mathcal{H}_\mathcal{K}$ is the subgraph consisting of vertices in the d-level and the $(d-1)$-level of of $\mathcal{H}_\mathcal{K}$ together with all the edges connecting these two levels.

The Hasse graph $\mathcal{H}_\mathcal{K}$ of a D-dimensional complex $\mathcal{K}$ has $(D+1)$ levels and D interfaces. To understand Morse theory in terms of Hasse graph, one needs to assign orientations to it.

Definition 2.1 (Oriented Hasse Graph, Up-edges, Down-edges)**.** If we assign orientations to all edges of Hasse graph $\mathcal{H}_\mathcal{K}$, we obtain an *oriented Hasse graph* denoted by $\overrightarrow{\mathcal{H}_\mathcal{K}}$. In graph $\overrightarrow{\mathcal{H}_\mathcal{K}}$, for any two simplices σ^d, τ^{d-1}, an edge $\tau^{d-1} \to \sigma^d \in \overrightarrow{\mathcal{H}_\mathcal{K}}$ going from lower dimensional simplex τ^{d-1} to a higher dimensional simplex σ^d is called an *up-edge*. The edge $\sigma^d \to \tau^{d-1} \in \overrightarrow{\mathcal{H}_\mathcal{K}}$ is called a *down-edge*.

If we orient $\mathcal{H}_\mathcal{K}$ in such a way that all edges are down-edges then this orientation corresponds to the trivial gradient vector field on complex $\mathcal{K}$, for which all simplices are critical. We call this the *default orientation* on $\mathcal{H}_\mathcal{K}$.

Matching based reorientation. Start with the default orientation on $\mathcal{H}_\mathcal{K}$. Associate a matching $\mathcal{M}$ to $\mathcal{H}_\mathcal{K}$. If an edge $\langle \tau^{d-1}, \sigma^d \rangle \in \mathcal{M}$ then *reverse* the orientation of that edge to $\tau^{d-1} \to \sigma^d \in \overrightarrow{\mathcal{H}_\mathcal{K}}$. We require the matching induced reorientation to be such that the graph $\overrightarrow{\mathcal{H}_\mathcal{K}}$ is a directed acyclic graph. Chari [7] first observed that every matching based orientation of $\mathcal{H}_\mathcal{K}$ that leaves the graph $\overrightarrow{\mathcal{H}_\mathcal{K}}$ acyclic corresponds to a unique gradient vector field on complex $\mathcal{K}$. For such a matching based acyclic orientation of the graph, every up-edge in the oriented Hasse graph corresponds to a gradient pair and every unmatched vertex corresponds to a critical simplex of the gradient vector field. Not every matching based orientation of $\mathcal{H}_\mathcal{K}$ will leave $\overrightarrow{\mathcal{H}_\mathcal{K}}$ acyclic. Figure 1 shows a simplicial complex and a matching based reorientation of the Hasse graph. We can now define the Max Morse Matching Problem more formally.

Definition 2.2 (Max Morse Matching Problem)**.** A discrete gradient vector field that maximizes the number of gradient pairs over the set of all discrete gradient vector fields on a simplicial complex $\mathcal{K}$ is known as a Maximum Morse Matching on $\mathcal{K}$. The Max Morse Matching Problem is to find such an optimal Morse Matching. In terms of Hasse graph, the Max Morse Matching Problem may be defined as the the maximum cardinality of an acyclic matching.

We now discuss a few properties of cycles and paths in a matching based orientation of $\mathcal{H}_\mathcal{K}$. Matching based orientations have the interesting property that all cycles are restricted to a fixed interface in the oriented Hasse graph. In other words, if a cycle were to span multiple interfaces in the Hasse graph, then it will violate the condition that the orientation is matching based. Similarly, all edges in a given path belong to a unique interface of the Hasse graph. Also, in a matching based orientation, source nodes and sink nodes in the d-interface are not involved in any cycles in the d-interface.

Definition 2.3 (Source and Sink Nodes)**.** A simplex σ^d is a *source node* for the d-interface if it has only outgoing edges to $d-1$ simplices. If in addition, simplex σ^d is matched to a $(d+1)$-simplex then it is as a *regular source node* for the d-interface, else it is a *critical source node*. Similarly, a simplex τ^{d-1} is a *sink node* for the d-interface if it has only incoming edges from d simplices. If τ^{d-1} is matched to a $(d-2)$-simplex then it is known as a *regular sink node* else it is known as a *critical sink node*.

[2]From here on, for the sake of brevity, while referring to the vertex in $\mathcal{H}_\mathcal{K}$ representing simplex σ^d, we drop the suffix $\mathcal{H}$ from $\sigma_\mathcal{H}^d$. i.e. Instead of referring to it as vertex $\sigma_\mathcal{H}^d$ we refer to it as simplex σ^d.

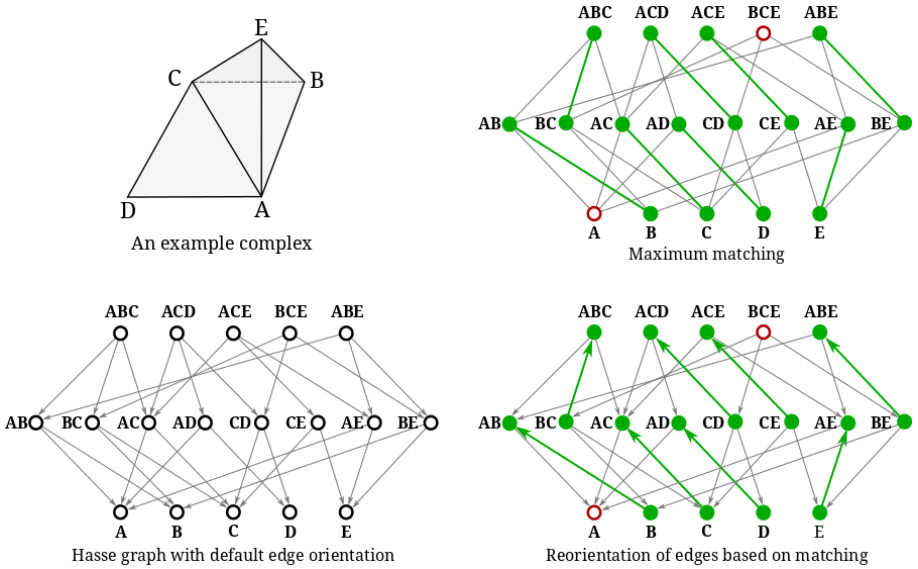

Figure 1: Consider a simplicial complex with five triangles (ABC, ACD, ACE, BCE, ABE). We obtain the oriented Hasse graph for a simplicial complex (left) and its matching induced orientation (right). Two simplices are critical (hollow) and others are regular (filled).

3 A $^{(D+1)}/_{(D^2+D+1)}$-factor approximation algorithm for simplicial complexes

We now describe an approximation algorithm for the Max Morse Matching Problem that is applicable to simplicial complexes. The idea is to first compute a maximum cardinality matching and, in a subsequent step, remove any cycles that maybe introduced due to the reorientation. The key steps are outlined in Algorithm 1. We begin with notes on notations and definitions.

Notation. When we denote an up-edge as $\chi(\alpha, \beta)$, we mean to say that it is an edge connecting simplex α^{d-1} to simplex β^d and is labelled as χ. We may write it either as $\chi(\alpha, \beta)$ or χ depending on whether we want to emphasize vertices incident on χ. The corresponding down-edge with reversed orientation is denoted as $\overline{\chi}$ or $\overline{\chi}(\beta, \alpha)$.

Definition 3.1 (Leading up-edges of an up-edge)**.** In an oriented Hasse graph $\overline{\mathcal{H}_\mathcal{K}}$, if we have an up-edge $\chi_1(\alpha_1, \beta_1)$ followed by a down-edge $\overline{\chi_2}(\beta_1, \alpha_2)$ followed by up-edge $\chi_3(\alpha_2, \beta_2)$ we say that χ_3 is a leading up-edge of χ_1.

Definition 3.2 (Facet-edges of a simplex)**.** In an oriented Hasse graph $\overline{\mathcal{H}_\mathcal{K}}$, for a simplex σ^d (where $d \geq 1$), the set of oriented edges between σ^d to $(d-1)$-simplices incident on σ^d (along with respective orientations) are known as the facet-edges of σ^d.

Given a Hasse graph $\mathcal{H}_\mathcal{K}$ on complex $\mathcal{K}$, Algorithm 1 begins by computing maximum cardinality graph matching on graph $\mathcal{H}_\mathcal{K}$ and then uses this matching to induce an orientation on $\mathcal{H}_\mathcal{K}$. Let $\overline{\mathcal{H}_\mathcal{K}}$ denote the oriented Hasse graph based on graph matching and $\mathcal{H}_\mathcal{V}$ denote the output graph. While there exists an up-edge χ in $\overline{\mathcal{H}_\mathcal{K}}$, we make χ a *seed-edge* and use it as a starting point for a BFS-like traversal on graph $\overline{\mathcal{H}_\mathcal{K}}$. This traversal is done using procedure **BFSComponent**() which returns a set of edges $\mathcal{C}_\chi$. The *edge-component* $\mathcal{C}_\chi$ of a seed edge χ is the set of edges discovered in the BFS-like traversal of graph $\overline{\mathcal{H}_\mathcal{K}}$, with χ as the start edge. Each time, we discover a new edge-component, we delete it from $\overline{\mathcal{H}_\mathcal{K}}$ and add it to $\mathcal{H}_\mathcal{V}$. We exit the while loop when all up-edges are exhausted.

If a simplex σ^d is either a critical node or a regular source node, then its facet-edges are not reachable in the BFS traversal through any of the up-edges in $\overline{\mathcal{H}_\mathcal{K}}$. In a final step, we include all remaining edges from $\overline{\mathcal{H}_\mathcal{K}}$ to $\mathcal{H}_\mathcal{V}$.

Algorithm 1 The Frontier Edges Algorithm

Input: Simplicial complex $\mathcal{K}$
Output: Graph $\mathcal{H}_{\mathcal{V}}$, an acyclic matching based orientation of Hasse graph $\mathcal{H}_{\mathcal{K}}$ of $\mathcal{K}$.

Construct Hasse graph $\mathcal{H}_{\mathcal{K}}$ of $\mathcal{K}$.
Perform maximum-cardinality graph matching on $\mathcal{H}_{\mathcal{K}}$.
Let $\overline{\mathcal{H}_{\mathcal{K}}}$ denote the matching induced reorientation of $\mathcal{H}_{\mathcal{K}}$ and $\mathcal{E}(\overline{\mathcal{H}_{\mathcal{K}}})$ its edge set.
Initialize the edge set of $\mathcal{H}_{\mathcal{V}}$, $\mathcal{E}(\mathcal{H}_{\mathcal{V}}) \leftarrow \emptyset$.
while $\exists \chi \in \mathcal{E}(\overline{\mathcal{H}_{\mathcal{K}}})$ such that χ is an up-edge **do**
  $\mathcal{C}_{\chi} \leftarrow$**BFSComponent**$(\overline{\mathcal{H}_{\mathcal{K}}}, \chi)$
  $\mathcal{E}(\overline{\mathcal{H}_{\mathcal{K}}}) \leftarrow \mathcal{E}(\overline{\mathcal{H}_{\mathcal{K}}}) \setminus \mathcal{C}_{\chi}$
  $\mathcal{E}(\mathcal{H}_{\mathcal{V}}) \leftarrow \mathcal{E}(\mathcal{H}_{\mathcal{V}}) \cup \mathcal{C}_{\chi}$
end while
$\mathcal{E}(\mathcal{H}_{\mathcal{V}}) \leftarrow \mathcal{E}(\mathcal{H}_{\mathcal{V}}) \cup \mathcal{E}(\overline{\mathcal{H}_{\mathcal{K}}})$
procedure BFSComponent$(\overline{\mathcal{H}_{\mathcal{K}}}, \chi)$
$\mathcal{C} \leftarrow \emptyset$
Initialize the queue $\mathcal{Q} \leftarrow \emptyset$
enqueue$(\mathcal{Q}, \chi)$
while $\mathcal{Q}$ is non-empty **do**
  $\chi_0(\alpha_0, \beta_0) \leftarrow$ **dequeue**$(\mathcal{Q})$
  $\mathcal{C} \leftarrow \mathcal{C} \cup$ **facetEdges**(β_0)
  for every leading up-edge $\chi_i(\alpha_i, \beta_i)$ of χ_0 **do**
    if the graph induced by edges in $(\mathcal{C} \cup$ **facetEdges**$(\beta_i))$ has cycles **then**
      Reverse orientation of χ_i in graph $\overline{\mathcal{H}_{\mathcal{K}}}$
      $\mathcal{C} \leftarrow \mathcal{C} \cup$ **facetEdges**(β_i)
    else
      enqueue$(\mathcal{Q}, \chi_i)$
    end if
  end for
end while
return $\mathcal{C}$
end procedure

The procedure **BFSComponent**() computes the component edges by processing edges from the queue one at a time. Let $\chi_0(\alpha_0, \beta_0)$ be the edge at the top of the queue. We add all the facet-edges of simplex β_0 to the edge-component $\mathcal{C}$. We now examine the leading up-edges of χ_0. If $\chi_i(\alpha_i, \beta_i)$ is a leading up-edge of χ_0 then we check if the addition of facet-edges of simplex β_i to $\mathcal{C}$ creates cycles. If it does then we classify χ_i as a *backward edge*, reverse the orientation of χ_i and add the facet-edges of β_i to $\mathcal{C}$. If this addition does not introduce cycles, then we classify χ_i as a *forward edge* and enqueue it in the queue of up-edges. Please refer to Figure 2. Enqueuing χ_i guarantees that at some stage when χ_i gets dequeued, we will end up adding facet-edges of simplex β_i to $\mathcal{C}$. When the queue is exhausted, $\mathcal{C}$ contains the entire edge-component of some seed-edge.

We first prove an acyclicity lemma on edge-components returned by procedure **BFSComponents**() in Algorithm 1**.**

Lemma 3.1. *The graph induced by edges in an edge-component is acyclic.*

Proof. Consider the graph induced by edges in edge-component $\mathcal{C}$ belonging to a d-interface. We know that an up-edge say χ_j is classified as a forward edge if and only if the inclusion of χ_j does not create a cycle with up-edges that were included prior to χ_j in edge-component $\mathcal{C}$. Hence, we can be sure that inclusion of set of all forward edges does not create cycles. Moreover, every time a backward edge, say $\chi_i(\alpha_i, \beta_i)$ is encountered, we include the inverse orientation of χ_i in $\mathcal{C}$ which creates a sink node at α_i and source node at β_i for the d-interface of the Hasse graph. Also, the $(d-1)$-simplices that were *visited* in a previous edge-component also act as sinks (since we restrict ourselves to edges induced by edge-component $\mathcal{C}$). Furthermore, every down-edge is incident on a $(d-1)$-simplex that is either a sink or a $(d-1)$-simplex incident on a forward

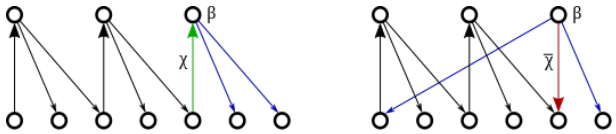

Figure 2: Two cases in BFSComponent(). Left: A forward edge χ is identified. The edge χ along with the down-edges incident on β are added to the edge-component. Right: A backward edge χ is identified. The edge $\bar{\chi}$ along with the down-edges incident on β are added to the edge-component.

edge. In either case, it is easy to see that all flow terminates at sinks making the graph induced by edges in a particular edge-component acyclic. □

Lemma 3.2. *The output graph $\mathcal{H}_\mathcal{V}$ is acyclic.*

Proof. We prove this claim via induction over sequential addition of edge-components.
Base Case: To begin with the output graph $\mathcal{H}_\mathcal{V}$ is the empty graph. From Lemma 3.1, we know that the graph induced by edges in an edge-component is acyclic. So $\mathcal{H}_\mathcal{V}$ remains acyclic following the addition of the first edge-component $\mathcal{C}_1$ to $\mathcal{H}_\mathcal{V}$.
Inductive Hypothesis: Suppose that following the addition of edges belonging to i^{th} edge-component $\mathcal{C}_i$, $\mathcal{H}_\mathcal{V}$ remains acyclic.
Now, we need to prove that following the addition of edges belonging to $\mathcal{C}_{i+1}$, $\mathcal{H}_\mathcal{V}$ remains acyclic. To begin with using Lemma 3.1, we note that the graph induced by $\mathcal{C}_{i+1}$ is acyclic. So, if there does exist a cycle in $\mathcal{H}_\mathcal{V}$ following the addition of $\mathcal{C}_{i+1}$, then a forward up-edge of this cycle must belong to $\mathcal{C}_{i+1}$ and a forward up-edge must belong to an edge-component $\mathcal{C}_{j_k}$ where $j_k < (i+1)$. In particular, this means that there exists a down-edge belonging to a component $\mathcal{C}_{j_k}$ that is incident on simplex α_1 such that a forward edge $\chi_1(\alpha_1, \beta_1) \in \mathcal{C}_{i+1}$. But, if α_1 was reachable while traversing $\mathcal{C}_{j_k}$ then $\chi_1(\alpha_1, \beta_1)$ would have been classified as a forward edge in $\mathcal{C}_{j_k}$ i.e. $\chi_1(\alpha_1, \beta_1) \in \mathcal{C}_{j_k}$ – a contradiction. Hence, such cycles do not exist. Finally, in line 10 of Algorithm 1, after having added all edge-components, we add all the facet-edges of d-simplices that are either unmatched or facet-edges of d-simplices that are matched to one of their cofacets. In such cases, they act as source nodes within d-interfaces and do not introduce cycles because all cycles are restricted to the d-interface. □

Lemma 3.3. *The output graph $\mathcal{H}_\mathcal{V}$ is a matching based acylic orientation of undirected Hasse graph of the complex $\mathcal{H}_\mathcal{K}$.*

Proof. We first prove that $\mathcal{H}_\mathcal{V}$ is an orientation of $\mathcal{H}_\mathcal{K}$ i.e. for every undirected edge in $\mathcal{H}_\mathcal{K}$ there is a corresponding directed edge in $\mathcal{H}_\mathcal{V}$. To prove this we will show that for every simplex β^d, all undirected edges from β to its facets in $\mathcal{H}_\mathcal{K}$ has a corresponding oriented edge in $\mathcal{H}_\mathcal{V}$.
Case 1: Suppose that β is matched to one of its facets in max-cardinality matching induced oriented graph $\overline{\mathcal{H}_\mathcal{K}}$. Then this up-edge incident on β was classified either as a forward edge or as a backward edge. In either case, all its facet-edges are inserted in $\mathcal{H}_\mathcal{V}$ in procedure **BFSComponent**().
Case 2: Now suppose that β is either unmatched or it is matched to one of its cofacets. Then clearly, none of its facet-edges can be reached through a graph traversal that starts with some up-edge in $\overline{\mathcal{H}_\mathcal{K}}$. Therefore, these facet-edges are not a part of any of the edge-components and they are all down-edges. However, in line 10 of Algorithm 1, all these *remainder* edges are included in $\mathcal{H}_\mathcal{V}$.
Since the above two cases hold true for every simplex σ^d with $d \geq 1$, this proves that $\mathcal{H}_\mathcal{V}$ is an orientation of graph $\mathcal{H}_\mathcal{K}$. Also, given the fact that the up-edges that are included are subset of edges coming from cardinality bipartite matching, clearly the orientation of $\mathcal{H}_\mathcal{V}$ is matching based. In Lemma 3.2, we already proved that graph $\mathcal{H}_\mathcal{V}$ is acyclic. Hence proved. □

Definition 3.3 (Classified Edges, Frontier Edges)**.** An edge marked within the **BFSComponent**() as forward or backward is called a *classified edge*. A leading up-edge that is not yet classified is called a *frontier edge*.

We establish the approximation ratio using a counting argument that works specifically for simplicial complexes. We refer to this argument as the frontier edges argument. The main idea involves a method

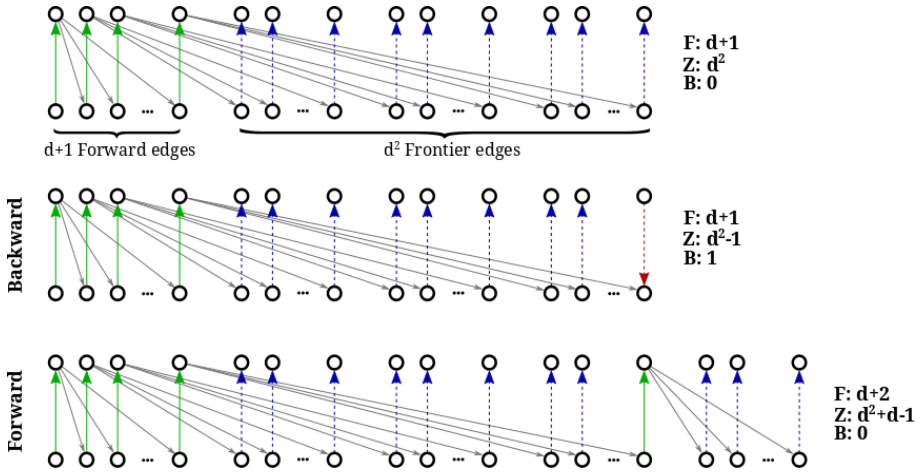

Figure 3: The frontier edges argument. Top: The base case with $d+1$ forward edges, d^2 frontier edges and no backward edges. The ratio of forward edges and total number of up-edges is $\frac{d+1}{d^2+d+1}$. Middle: When one of the frontier edges is classified as backward, the ratio remains the same. Bottom: When one of the frontier edges is classified as a forward edge, the ratio improves to $\frac{d+2}{d^2+2d+1}$.

of counting that we describe now. Suppose we are processing an edge-component that belongs to the d-interface of the Hasse graph for some $d \leq D$. Let the iterator variable i count the number of up-edges in the edge-component that have so far been classified as either forward or backward. Suppose at the end of i^{th} iteration, there are $|\mathcal{F}_i|$ number of forward edges, $|\mathcal{B}_i|$ number of backward edges and $|\mathcal{Z}_i|$ number of frontier edges, then our approximation ratio is calculated as ${}^{|\mathcal{F}_i|}/_{(|\mathcal{F}_i|+|\mathcal{B}_i|+|\mathcal{Z}_i|)}$. In other words, we assume the worst case scenario where all the frontier edges are *possibly backward*. In every iteration of the BFS, we classify one of the frontier edges as a forward edge or a backward edge and then update the ratio until we exhaust the entire edge-component. In the $(i+1)^{th}$ iteration, if a frontier edge is classified as a forward edge then the number of forward edges will be $|\mathcal{F}_{i+1}| = (|\mathcal{F}_i|+1)$ and the number of frontier edges will be $|\mathcal{Z}_{i+1}| = (|\mathcal{Z}_i|+d-1)$. If a frontier edge is classified as a backward edge then the number of backward edges will be $|\mathcal{B}_{i+1}| = (|\mathcal{B}_i|+1)$ and the number of frontier edges will be $|\mathcal{Z}_{i+1}| = (|\mathcal{Z}_i|-1)$.

Lemma 3.4. *The number of forward edges in an edge-component belonging to the d-interface of the Hasse graph is at least* ${}^{(d+1)}/_{(d^2+d+1)}$ *fraction of the total number of up-edges in the edge-component.*

Proof. We will use induction to prove our claim.

Base Case: The seed edge χ_0 of the edge-component is naturally a forward edge. We note that any cycle in the Hasse graph of a simplicial complex has minimum length 6 and involves at least 3 up-edges. Since this does not hold for general regular cell complexes, simplicial input is crucial for the proof to work. Cycles do not appear until after two iterations. These two iterations constitute the base case. Therefore, $|\mathcal{F}_1| = 1$, $|\mathcal{B}_1| = 0$ and $|\mathcal{Z}_1| = 0$. Also, the leading up-edges of χ_0 are also forward edges. If χ_0 has no leading up-edges then the edge-component is exhausted and ${}^{|\mathcal{F}_1|}/_{(|\mathcal{F}_1|+|\mathcal{B}_1|)} = 1$. If χ_0 has K leading up-edges, each such edge has, in turn, at most j_k leading up-edges then the total number of forward edges will be $|\mathcal{F}_2| = 1+K$, $|\mathcal{B}_2| = 0$ and $|\mathcal{Z}_2| = \sum_{k=1}^{K} j_k$. It is easy to check that the worst case for ratio ${}^{|\mathcal{F}_2|}/_{(|\mathcal{F}_2|+|\mathcal{B}_2|+|\mathcal{Z}_2|)}$ occurs when $K = d$ and $j_k = d$ for each k. This gives us the worst case ratio for the quantity ${}^{|\mathcal{F}_2|}/_{(|\mathcal{F}_2|+|\mathcal{B}_2|+|\mathcal{Z}_2|)}$ to be ${}^{(d+1)}/_{(d^2+d+1)}$. Please refer to Figure 3

Induction Step: Our induction hypothesis says that after i iterations of BFS, the ratio ${}^{|\mathcal{F}_i|}/_{(|\mathcal{F}_i|+|\mathcal{B}_i|+|\mathcal{Z}_i|)} \geq {}^{(d+1)}/_{(d^2+d+1)}$. For the $(i+1)^{th}$ iteration, suppose one of the frontier edges is classified as a forward edge. Then $|\mathcal{F}_{i+1}| = (|\mathcal{F}_i|+1)$ and $|\mathcal{Z}_{i+1}| \leq (|\mathcal{Z}_i|+d-1)$. Note that $(|\mathcal{Z}_i|+d-1)$ is the worst case estimate for $|\mathcal{Z}_{i+1}|$ assuming that the newly included forward edge has d leading up-edges. Therefore, the numerator of the ratio ${}^{|\mathcal{F}_i|}/_{(|\mathcal{F}_i|+|\mathcal{B}_i|+|\mathcal{Z}_i|)}$ increases by 1 whereas the denominator increases by d. Also we have ${}^{1}/_{d} > {}^{(d+1)}/_{(d^2+d+1)}$. Using the elementary fact that if $\frac{A}{B} \geq \frac{E}{F}$ and $\frac{C}{D} \geq \frac{E}{F}$ then $\frac{(A+C)}{(B+D)} > \frac{E}{F}$ for non-negative values of A, B, C, D, E

and F, we get:

$$\begin{aligned}\frac{|\mathcal{F}_{i+1}|}{(|\mathcal{F}_{i+1}|+|\mathcal{B}_{i+1}|+|\mathcal{Z}_{i+1}|)} &\geq \frac{|\mathcal{F}_i|+1}{(|\mathcal{F}_i|+|\mathcal{B}_i|+|\mathcal{Z}_i|)+d} \\ &\geq \frac{(d+1)+1}{(d^2+d+1)+d} \\ &> \frac{(d+1)}{(d^2+d+1)}\end{aligned}$$

On the other hand if a frontier edge is classified as a backward edge then $|\mathcal{B}_{i+1}| = (|\mathcal{B}_i|+1)$ and $|\mathcal{Z}_{i+1}| = (|\mathcal{Z}_i|-1)$. So, the numerator and the denominator of the ratio ${}^{|\mathcal{F}_{i+1}|}/_{(|\mathcal{F}_{i+1}|+|\mathcal{B}_{i+1}|+|\mathcal{Z}_{i+1}|)}$ remain unchanged which gives us ${}^{|\mathcal{F}_{i+1}|}/_{(|\mathcal{F}_{i+1}|+|\mathcal{B}_{i+1}|+|\mathcal{Z}_{i+1}|)} = {}^{|\mathcal{F}_i|}/_{(|\mathcal{F}_i|+|\mathcal{B}_i|+|\mathcal{Z}_i|)}$. In both cases, the bound holds after $(i+1)$ iterations. □

Since every edge-component that belongs to a d-interface achieves a ratio of at least ${}^{(d+1)}/_{(d^2+d+1)}$ edges, if we sum over all the edge-components we get the ratio ${}^{(d+1)}/_{(d^2+d+1)}$ for that d-interface. In other words, we preserve at least ${}^{(d+1)}/_{(d^2+d+1)}$ of the total number of matchings at every d-interface. The ratio ${}^{(d+1)}/_{(d^2+d+1)}$ becomes worse with increasing d. So the worst case ratio is ${}^{(D+1)}/_{(D^2+D+1)}$ where D is the dimension of the complex. Therefore, we get the following result on the approximation ratio.

Theorem 3.5. *Algorithm 1 computes a ${}^{(D+1)}/_{(D^2+D+1)}$-factor approximation for Max Morse Matching Problem on simplicial complexes of dimension D.*

Proof. Let $|\mathcal{M}|$ denote the cardinality of maximum matching. Note that $2|\mathcal{M}|$ is an upper bound on Max Morse Matching i.e optimal number of regular simplices $\leq 2|\mathcal{M}|$. Since we preserve at least ${}^{(D+1)}/_{(D^2+D+1)}$ of these matchings, the number of regular simplices we obtain is at least $2\frac{(D+1)}{(D^2+D+1)}|\mathcal{M}| \geq \frac{(D+1)}{(D^2+D+1)}OPT$. Therefore, Algorithm 1 provides a $\frac{D+1}{(D^2+D+1)}$-factor approximation for the Max Morse Matching Problem on simplicial complexes. □

3.1 A ${}^{5}/_{11}$-factor Approximation for 2-dimensional simplicial complexes using Frontier Edges Algorithm

In this section, we observe that we can further tighten our analysis of Algorithm 1 by restricting the problem to 2-dimensional simplicial complexes. We exploit the geometry of 2-complexes as proved in Lemma 3.6 in order to establish an improved ratio in the base case.

Lemma 3.6. *If α is a forward edge and β_1 is a leading forward edge of edge α and if γ_1 and γ_2 are leading up-edges of β_1 then only one of the two edges γ_1 and γ_2 can possibly be a backward edge that creates a cycle with edge α.*

Proof. Without loss of generality, in this proof, we will use concrete labeling of simplices. We make an elementary geometric observation to prove this claim. Suppose α is a forward edge between a 1-simplex say AB matched to a 2-simplex ABC. So α can alternatively be denoted as edge AB-ABC. Now suppose 1-simplex BC is matched to another 2-simplex BCD constituting forward edge β_1, then of the two 1-simplices BD and CD, BD can possibly match a 2-simplex say BDA which effectively makes edge BD-BDA (say γ_1) a backward edge. However it is impossible to have a forward edge incident on 1-simplex CD (say γ_2) that is also simultaneously incident on 1-simplex AB since any 2-simplex has at most three vertices. Hence proved. □

Lemma 3.7. *The number of forward edges is at least ${}^{5}/_{11}$ fraction of the total number of up-edges in the edge-component.*

Proof. Once again, we will use induction to prove our claim.
Base Case: In case of 2-manifolds, we can count up to three levels of BFS for base case, which in turn gives us an improvement in ratio. The seed edge α of the edge-component is evidently a forward edge. We note that any cycle in the Hasse graph of a simplicial complex has minimum length 3. Therefore, $|\mathcal{F}_1| = 1$, $|\mathcal{B}_1| = 0$ and $|\mathcal{Z}_1| = 0$. Also, the leading up-edges of α (if any) are also forward edges. If α has no leading

up-edges then the edge-component is exhausted and $|\mathcal{F}_1|/(|\mathcal{F}_1|+|\mathcal{B}_1|) = 1$. If α has one leading up-edge β_1, then $|\mathcal{F}_2| = 2$, $|\mathcal{B}_2| = 0$ and $|\mathcal{Z}_2| = 2$. Therefore, $|\mathcal{F}_2|/(|\mathcal{F}_2|+|\mathcal{B}_2|+|\mathcal{Z}_2|) = 1/2$. If α has two leading up-edges β_1 and β_2, then $|\mathcal{F}_2| = 3$, $|\mathcal{B}_2| = 0$ and $|\mathcal{Z}_2| = 4$. Therefore, $|\mathcal{F}_2|/(|\mathcal{F}_2|+|\mathcal{B}_2|+|\mathcal{Z}_2|) = 3/7$. By Lemma 3.6, both leading up-edges of β_1, γ_1 and γ_2, can not be backward. So suppose one of them (say γ_1) is backward and γ_2 is forward then $|\mathcal{F}_3| = 4$, $|\mathcal{B}_3| = 1$ and $|\mathcal{Z}_3| = 4$ and therefore $|\mathcal{F}_3|/(|\mathcal{F}_3|+|\mathcal{B}_3|+|\mathcal{Z}_3|) = 4/9$. Similarly, we must also consider the leading up-edges of β_2 of which at most one of them can be backward. The worst case occurs for the configuration when exactly one leading up-edge each of β_1 and β_2 are backward. This configuration gives $|\mathcal{F}_3| = 5$, $|\mathcal{B}_3| = 2$ and $|\mathcal{Z}_3| = 4$ and hence $|\mathcal{F}_3|/(|\mathcal{F}_3|+|\mathcal{B}_3|+|\mathcal{Z}_3|) = 5/11$.
Induction Step: Our induction hypothesis is that following i iterations of BFS, the ratio $|\mathcal{F}_i|/(|\mathcal{F}_i|+|\mathcal{B}_i|+|\mathcal{Z}_i|) \geq 5/11$. For the $(i+1)^{th}$ iteration, suppose one of the frontier edges is classified as a forward edge. Then $|\mathcal{F}_{i+1}| = (|\mathcal{F}_i|+1)$ and $|\mathcal{Z}_{i+1}| = (|\mathcal{Z}_i|+1)$. Therefore, the numerator of the ratio $|\mathcal{F}_i|/(|\mathcal{F}_i|+|\mathcal{B}_i|+|\mathcal{Z}_i|)$ increases by 1 whereas the denominator increases by 2. However, since $1/2 > 5/11$, we have $|\mathcal{F}_{i+1}|/(|\mathcal{F}_{i+1}|+|\mathcal{B}_{i+1}|+|\mathcal{Z}_{i+1}|) = (5+1)/(11+1) > 5/11$. On the other hand if a frontier edge is classified as a backward edge then $|\mathcal{B}_{i+1}| = (|\mathcal{B}_i|+1)$ and $|\mathcal{Z}_{i+1}| = (|\mathcal{Z}_i|-1)$. So numerator and denominator of ratio $|\mathcal{F}_{i+1}|/(|\mathcal{F}_{i+1}|+|\mathcal{B}_{i+1}|+|\mathcal{Z}_{i+1}|)$ remain unchanged which gives us $|\mathcal{F}_{i+1}|/(|\mathcal{F}_{i+1}|+|\mathcal{B}_{i+1}|+|\mathcal{Z}_{i+1}|) = |\mathcal{F}_i|/(|\mathcal{F}_i|+|\mathcal{B}_i|+|\mathcal{Z}_i|)$. When all up-edges of the edge-component are exhausted, we don't have anymore frontier edges and the ratio for the edge-component after processing $|\mathcal{F}|$ forward edges and $|\mathcal{B}|$ backward edges will be $|\mathcal{F}|/(|\mathcal{F}|+|\mathcal{B}|)$ and by our inductive argument the ratio will be at least $5/11$. □

Once again since every edge-component achieves a ratio of at least $5/11$ edges, if we sum over all the edge-components we get the following theorem as an immediate outcome of the lemma above.

Theorem 3.8. *Algorithm 1 is a $5/11$-factor approximation algorithm for Max Morse Matching Problem when restricted to 2-dimensional simplicial complexes.*

4 Approximation algorithms for simplicial manifolds

4.1 A $2/(D+1)$-factor approximation algorithm for simplicial manifolds

We will restrict our attention to manifolds without boundary. The key idea in Algorithm 2 is that the matching is constructed within one d-interface at a time, starting with the lowest interface and ending with the highest one. For manifolds, this is advantageous because it allows us to count matched/critical simplices differently. In particular, every d-simplex (where $1 \leq d \leq D-1$), is given two chances to get matched. Please refer to Figure 4. We first try to match a d-simplex say σ^d, while constructing the Morse matching for the d-interface. If σ^d remains critical for the d-interface then we try to match it for the $(d+1)$-interface. The trick of giving a second chance to critical simplices works fine for all dimensions except for D-dimensional critical simplices. Fortunately, for manifolds, we can easily design a vector field with only one critical simplex for dimension D. Since non-manifold-complexes may have unbounded number of critical D-simplices the analysis becomes non-trivial. For Algorithm 2, one may still derive approximation bounds for non-manifold complexes by using a line of reasoning analogous to one used in Section 4.2.1.

Algorithm 2 , exploits special structures at the lowest and highest interface. For instance, for any D-dimensional manifold, there are well known algorithms in literature [4, 19, 15] for designing optimal gradient vector field for the 1-interface and the D-interface. See Appendix A of [4]. As noted in [4], we can associate a special graph structure to the D-interface.

Definition 4.1 (Dual Graph)**.** The dual graph of a simplicial D-dimensional manifold $\mathcal{K}$, denoted by $\Gamma(\mathcal{K})$ is the graph whose vertices represent the D-simplices of $\mathcal{K}$ and whose edges join two D-simplices with a common $(D-1)$-facet.

For the sake of completeness, we describe the optimal algorithms for 1-interface and the D-interface in Appendix A.

Like in Algorithm 1, we first obtain the Hasse graph of complex $\mathcal{K}$. We extract the d-interface of the Hasse graph. For the d-interface, we design a Morse matching and reorient the output graph $\mathcal{H}_\mathcal{V}$ based on it. We then delete all the regular simplices of the d-interface and the critical $(d-1)$-simplices. This updated Hasse graph is available for the next iteration when the $(d+1)$-interface is extracted and so on.

Algorithm 2 The Interface Algorithm

Input: Simplicial complex $\mathcal{K}$

Output: Graph $\mathcal{H}_{\mathcal{V}}$, an acyclic matching based orientation of Hasse graph $\mathcal{H}_{\mathcal{K}}$ of $\mathcal{K}$.

▷ **Notation:** [$\mathcal{C}_d^{d-1}$ denotes the critical $(d-1)$-simplices for d-interface. $\mathcal{R}_d$ is the set of all regular simplices for d-interface and $\mathcal{M}_d$ is the set of gradient pairs for d-interface. $\mathcal{E}(\mathcal{H}_{\mathcal{V}})$ denotes the edge set of $\mathcal{H}_{\mathcal{V}}$.]
Construct Hasse graph $\mathcal{H}_{\mathcal{K}}$ of $\mathcal{K}$.
$\mathcal{E}(\mathcal{H}_{\mathcal{V}})$ is initialized to default down-edge orientation on all edges.
for $d = 1$ to D **do**
  $\mathcal{G}^d \leftarrow$ **extractdInterface**$(\mathcal{H}_{\mathcal{K}}, d)$
  if $d = 1$ **then** Apply **1ComplexOpt**$(\mathcal{G}^1)$
  else if $d = D$ **then** Apply **manifoldOpt**$(\mathcal{G}^D)$
  else Apply **intermediateApx**$(\mathcal{G}^d, d)$
  end if
end for

procedure DELETEANDREORIENT$(\mathcal{C}_d^{d-1}, \mathcal{R}_d, \mathcal{M}_d)$
Reorient edges of $\mathcal{H}_{\mathcal{V}}$ based on matchings in edge set $\mathcal{M}_d$ for the d-interface)
Delete nodes $\{\mathcal{C}_d^{d-1}, \mathcal{R}_d\}$ from $\mathcal{H}_{\mathcal{K}}$
end procedure

procedure 1COMPLEXOPT$(\mathcal{W})$
Apply the optimal algorithm on $\mathcal{W}$. (See **DFSoptimal**() in Appendix A).
deleteAndReorient$(\mathcal{C}_1^0, \mathcal{R}_1, \mathcal{M}_1)$
end procedure

procedure MANIFOLDOPT$(\mathcal{W})$
Apply the optimal algorithm on the dual graph. (See **DFSoptimal**() in Appendix A).
Reorient edges of $\mathcal{H}_{\mathcal{V}}$ based on matchings in edge set $\mathcal{M}_D$ for the D-interface.
end procedure

procedure INTERMEDIATEAPX$(\mathcal{W}, d)$
Apply Algorithm 1 described in Section 3 on graph $\mathcal{W}$.
For every unmatched simplex τ^{d-1} such that all its cofacets $\sigma_1^d \dots \sigma_K^d$ are also unmatched, choose one of the simplices $\sigma_i^d, i \in [1, K]$ and introduce the matching $\langle \tau, \sigma_i \rangle$.
deleteAndReorient$(\mathcal{C}_d^{d-1}, \mathcal{R}_d, \mathcal{M}_d)$
end procedure

For the 1-interface and the D-interface, the optimal algorithms are applied to design the Morse matchings. For d-interfaces where $1 < d < D$, procedure **intermediateApx**() of Algorithm 2 is applied to design the Morse matchings.

We now describe procedure **intermediateApx**() for designing gradient vector field on the d-interface $\mathcal{G}^d$**.** Algorithm 1 is essentially a maximum-matching followed by BFS-style cycle removal and hence can be performed on any bipartite graph. In particular, we apply it on graph $\mathcal{G}^d$ for $1 < d < D$. After cycle removal (from Algorithm 1) we may have a situation where we have an unmatched simplex τ such that all its cofacets are also unmatched. In that case, we match τ with one of its cofacets. We perform this operation for all unmatched $(d-1)$-simplices whose cofacets are also unmatched. This completes Morse matching for the d-interface. In procedure **deleteAndReorient**(), if σ^d is incident on simplex τ^{d-1} and if τ is regular at the $(d-1)$-interface then we are justified in deleting it while processing the d-interface, since τ is a regular sink node for d-interface. The deletion of critical nodes does not affect the behavior of Algorithm 2 per se. We delete them here because the procedure **intermediateApx**() is used as a subroutine in Algorithm 3 where this deletion is crucial.

Lemma 4.1. *The orientation of $\mathcal{G}^d$ as computed by Algorithm 2 is acyclic.*

Proof. Algorithm 1 provides an acyclic matching-based orientation of d-interface $\mathcal{G}^d$. So step 1 of **intermediateApx**() does not introduce any cycles. Now consider an unmatched simplex τ^{d-1} such that all its cofacets $\sigma_1^d \dots \sigma_K^d$ are also unmatched. For a directed acyclic graph there is an ordering relation $\alpha > \beta$ if

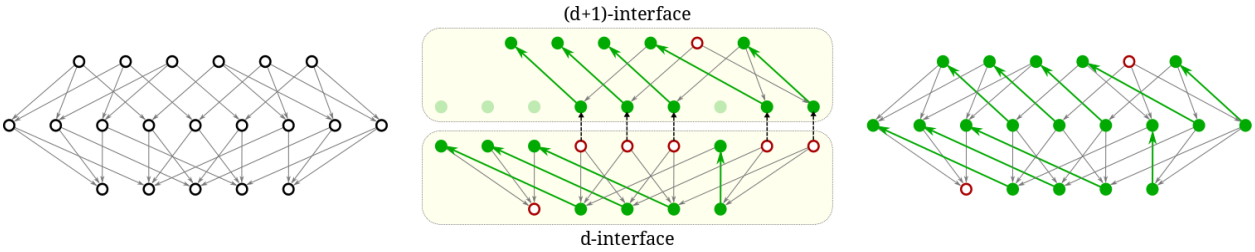

Figure 4: An example illustrating Algorithm 2. Left: The input Hasse graph. Middle: The d-simplices matched during execution of the algorithm on the d-interface are deleted from the $(d+1)$-interface. Right: The final matching is obtained by combining matchings from all interfaces.

there is a directed path from vertex α to vertex β. Clearly there is no ordering relation among $\sigma_1^d \dots \sigma_K^d$ since they are all critical. Introduction of the matching $\langle \tau, \sigma_i \rangle$ introduces the ordering relations of the type $\sigma_j > \sigma_i$ for all $j \in [1, K]$ and $j \neq i$. Therefore matching introduced by step 2 does not introduce any cycles and hence the orientation of $\mathcal{G}^d$ as computed by Algorithm 2 is acyclic. □

Lemma 4.2. *The orientation of the output graph $\mathcal{H}_\mathcal{V}$ is acyclic.*

Proof. From Lemma 4.1, we conclude that the orientation for every d-interface $\mathcal{G}^d$ where $1 < d < D$ is acyclic. Further, optimal acyclic matchings are computed for 1-interface and D-interface respectively. Combining these two facts and along with the observation that every directed path is restricted to a unique d-interface, we conclude that the orientation of output graph $\mathcal{H}_\mathcal{V}$ is acyclic. □

Now we introduce an idea that will help us prove approximation bounds for Algorithm 2. For the d-interface $\mathcal{G}_d$, let τ^{d-1} be a critical simplex and let the set of cofacets of τ that are regular be $\{\beta_1, \beta_2, \dots, \beta_K\}$. From line 25 of procedure **intermediateApx**(), we know that this set is non-empty. Let β_i where $i \in [1 \dots K]$ be a cofacet of τ with minimum index after performing a topological sort on the d-interface[3]. Now let α_i be such that $\alpha_i \prec \beta_i$ and $\langle \alpha_i, \beta_i \rangle$ is a gradient pair. Then we can associate a *canonical triplet* $\langle \langle \alpha_i, \beta_i \rangle, \tau \rangle$ to critical simplex τ^{d-1}. Note that such a unique canonical triplet is associated to every critical $(d-1)$-simplex.

Lemma 4.3. *Algorithm 2 computes a $2/(d+2)$-factor approximation to the Max Morse Matching restricted to the d-interface, $1 < d < D$, of the Hasse graph of the D-dimensional manifold.*

Proof. Let $\left\langle \alpha_i^{d-1}, \beta_i^d \right\rangle$ be a gradient pair. β_i^d has $(d+1)$ facets of which at least one (namely α_i) is regular. Therefore, the gradient pair $\langle \alpha_i, \beta_i \rangle$ appears in at most d canonical triplets. We group $\langle \alpha_i, \beta_i \rangle$ with all the critical $(d-1)$-simplices that contain $\left\langle \alpha_i^{d-1}, \beta_i^d \right\rangle$ in their canonical triplets. Each critical $(d-1)$-simplex appears in a unique group. Each group contains at least two regular simplices and at most d-critical simplices. So for every group we have the following ratio

$$\frac{\#\text{matched simplices}}{\#\text{total simplices}} \geq \frac{2}{(d+2)}$$

Hence, we obtain the approximation ratio of $2/(d+2)$ for the d-interface. □

The minimum of the ratio $2/(d+2)$ over all d, $1 < d < D$ is $2/(D+1)$. The 1-interface contributes to a single critical simplex when the optimal algorithm is employed (See Appendix A).

Finally, we consider the D-interface in the lemma below.

Lemma 4.4. *At least one $(D-1)$-simplex is matched at the conclusion of construction of Morse matching at the $(D-1)$-interface.*

Proof. Consider an arbitrary $(D-1)$-simplex α_1. Let $\partial\alpha_1$ denote the simplical boundary of an arbitrary $(D-1)$-simplex α_1. Clearly, $\partial\alpha_1$ is a $(D-2)$-manifold without boundary. Therefore, using Morse inequalities there exists at least one $(D-2)$-simplex in $\partial\alpha$ that remains unmatched upon construction of Morse matching

[3]We do not actually perform topological sort on the d-interface, but need it for making an argument.

at the $(D-2)$-interface. When we look at boundaries of all such $(D-1)$-simplices α_i, we may find several unmatched $(D-2)$-simplices upon conclusion of Morse matching at $(D-2)$-interface. Since one or more $(D-2)$-simplices remain unmatched at the start of the construction of Morse matching at the $(D-1)$-interface, at least one $(D-2)$-simplex can be matched to a $(D-1)$-simplex (without introducing cycles). □

Lemma 4.5. *After constructing Morse matching at the D-interface, the following ratio holds true:*

$$\frac{\#\text{matched simplices}}{\#\text{total simplices}} \geq \frac{4}{D+3}$$

Proof. Let k denote the number of $(D-1)$-simplices that are matched at the conclusion of construction of Morse matching at the $(D-1)$-interface. Using Lemma 4.4, we have $k \geq 1$.
Now consider the dual graph structure of the D-interface. The vertex degree of the dual graph is bounded by $D+1$. So the total number of edges in the dual graph is smaller than $(D+1)/2$. Applying the optimal algorithm in Appendix A, ensures that we have only one critical simplex in the dual graph. If N is the number of vertices in the dual graph, the following ratio holds true:

$$\frac{\#\text{matched simplices}}{\#\text{total simplices}} \geq \frac{2(N-1)}{(\frac{D+1}{2}+1)(N-1)+1-k} \geq \frac{2(N-1)}{(\frac{D+1}{2}+1)(N-1)} \geq \frac{4}{D+3}$$

□

Note that $4/(D+3) > 2/(D+1)$ for all $D \geq 3$. So the worst ratio over all d-interfaces, where $1 \leq d \leq D$, is $\frac{2}{(D+1)}$. Since the optimal number of regular simplices is bounded by the total number of simplices, we get the following theorem.

Theorem 4.6. *For $D \geq 3$, Algorithm 2 provides a $2/(D+1)$-factor approximation for the Max Morse Matching problem for manifolds without boundary.*

We will like to make two remarks here regarding the approximation factor. Firstly, the ratio is not affected by line 24 (first step) of procedure **intermediateApx**(). It depends entirely on line 25 (second step) of **intermediateApx**(). We include a matching based preprocessing step prior to applying the second step because in practice, doing so, gives significantly better results. Secondly, the approximation ratio is over the *total number of simplices*. In that sense, Algorithm 2 and its analysis helps further our understanding of combinatorial construction of manifolds. In other words, irrespective of the complex size, the homology or the presence of non-collapsible elements, we can always collapse at least $2/(D+1)$ number of simplices in that manifold!

4.2 A $2/D$-factor approximation algorithm for simplicial manifolds

Once again we restrict our attention to simplicial manifolds without boundary. We build on Algorithm 2 by exploiting a finer substructure within each interface to obtain a further improvement in ratio for simplicial manifolds. We begin with some definitions.

Definition 4.2 (facet degree, min-facet simplex of the d-interface)**.** The number of facets incident on a simplex is defined as its *facet degree*. For the d-interface, consider the subset of d-simplices $\mathcal{S}$ with at least one facet. We say that a d-simplex is a *min-facet simplex* if over all simplices in $\mathcal{S}$, it has the minimum number of facets.

Definition 4.3 (min-facet component of the d-interface)**.** A *min-facet component* is a subgraph of the d-interface that is a maximal connected graph induced by a set of min-facet simplices of the d-interface.

Like in Algorithm 2, we process the Hasse graph one d-interface at a time starting with the 1-interface and terminating with the D-interface. Also, like in Algorithm 2, we use optimal algorithms to process the 1-interface and the D-interface of the Hasse graph. Only the intermediate interfaces are processed differently. The procedure for handling intermediate interfaces is outlined in Algorithm 3.

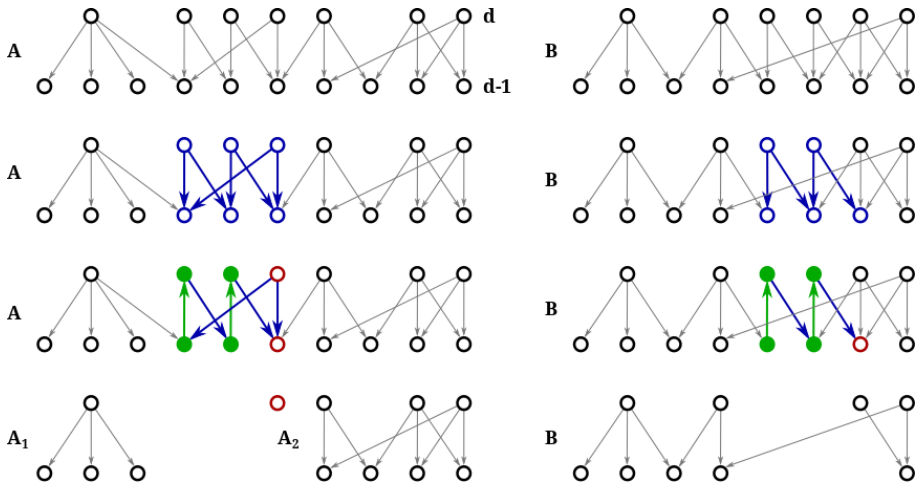

Figure 5: Algorithm 3 processes the min-facet component in the d-interface (bold edges). Regular simplices are denoted by filled vertices. Critical simplices and unprocessed simplices are denoted by hollow vertices. Left: Deletion of $(d-1)$-simplices of a min-facet component disconnects the graph. Right: Deletion of $(d-1)$-simplices of the new min-facet component keeps the graph connected. This process continues until none of the d-simplices have any facets left.

Algorithm 3 The Min-Facet Component Algorithm

procedure INTERAPXMINFACET($\mathcal{G}^d$)
  while sizeOfMinFacet($\mathcal{G}^d$) > 0 **do**
    $\mathcal{F}_\mathcal{C} \leftarrow$ **extractMinFacetComponent**($\mathcal{G}^d$)
    Apply **intermediateApx**($\mathcal{F}_\mathcal{C}, d$) from Algorithm 2
  end while
end procedure

By design, procedure **intermediateApx**() from Algorithm 2 need not process the entire d-interface $\mathcal{G}^d$ at one go. It may take any subgraph of the d-interface as its input. The key idea is to iteratively compute Morse matching by executing **intermediateApx**() on a min-facet component and after designing a vector field on this component, we subsequently delete it from the d-interface $\mathcal{G}^d$. As a consequence, $\mathcal{G}^d$ grows increasingly sparse and when the entire d-interface has no edges left the while loop terminates. Figure 5 illustrates sample executions of the Algorithm 3.

Lemma 4.7. *If the d-interface of the Hasse graph is connected then there exists a gradient path connecting any two simplices α^{d-1} and β^{d-1}.*

Proof. Since the d-interface is connected there exists a path in the d-interface graph connecting any two vertices α^{d-1} and β^{d-1}. However, it is easy to see that any path connecting α^{d-1} and β^{d-1} in the d-interface graph is a gradient path in the discrete Morse theory sense. □

Lemma 4.8. *If the d-interface of the Hasse graph is connected then the $(d-1)$-interface is connected.*

Proof. Suppose the d-interface of the Hasse graph is connected. If any two arbitrary $(d-2)$-simplices γ_s and γ_d can be shown to be connected, then the $(d-1)$-interface is connected. To begin with let α_0 be any $(d-1)$-simplex with γ_s as its facet and α_k be any $(d-1)$-simplex with γ_d as its facet. If $\alpha_0 = \alpha_k$, there is nothing to prove. So for the remainder of the proof we shall assume that $\alpha_0 \neq \alpha_k$. Since the d-interface is connected, using Lemma 4.7, there exists a gradient path $\alpha_0, \beta_1, \alpha_1 \ldots, \alpha_{(k-1)}, \beta_k, \alpha_k$ connecting α_0 and α_k where all $\beta_i, i \in [1,k]$ are d-simplices and all $\alpha_i, i \in [0,k]$ are $(d-1)$-simplices. Now, since every d-simplex $\beta_i, i \in [1,k]$ is common to two $(d-1)$-simplices $\alpha_{(i-1)}$ and α_i belonging to the gradient path connecting α_0 and α_k, we know that $\alpha_{(i-1)}$ and α_i will share a facet which we denote by $\gamma_i^{(d-2)}$. In other words, we construct a new simplicial sequence $\mathcal{S} = \alpha_0, \gamma_1, \alpha_1 \ldots, \alpha_{(k-1)}, \gamma_k, \alpha_k$ from gradient path $\alpha_0, \beta_1, \alpha_1 \ldots, \alpha_{(k-1)}, \beta_k, \alpha_k$, where $\gamma_i = \alpha_{(i-1)} \cap \alpha_i$. However, note that in this case γ_i may possibly be equal to γ_j for some $i \neq j$. See Fig.6 for an example. Without loss of generality assume $\gamma_s \neq \gamma_1$ and $\gamma_d \neq \gamma_k$. We prove connectivity of γ_s and γ_d by induction. For base case, we note that γ_s is connected to γ_1 since γ_1 and γ_s are facets of simplex α_0. For induction step, suppose γ_s is connected to γ_i. Now consider the next two elements in sequence $\mathcal{S}$ namely α_i and $\gamma_{(i+1)}$. If $\gamma_i = \gamma_{(i+1)}$, then α_i makes no contribution towards finding a path connecting γ_s and γ_d and hence we ignore it. Else if $\gamma_i \neq \gamma_{(i+1)}$, then both γ_i and $\gamma_{(i+1)}$ are facets of α_i and hence γ_i is connected to $\gamma_{(i+1)}$ in the $(d-1)$ interface. By transitivity, γ_s is connected to $\gamma_{(i+1)}$, which completes the induction step. Finally, both γ_k and γ_d are facets of α_k and hence γ_k is connected to γ_d. By transitivity γ_s is connected to γ_d. This proves that there exists a subgraph of the $(d-1)$-interface that connects any two arbitrary $(d-2)$-simplices γ_s and γ_d. Hence proved. □

Lemma 4.9. *For a connected D-manifold without boundary, all d-interfaces are connected.*

Proof. Let K be the number of connected components of the D-interface. Then clearly $\beta_D = K$. Since a connected manifold without boundary has $\beta_D = 1$, we conclude that the D-interface is connected. Combining this fact with Lemma 4.8 implies that all d-interfaces are connected. □

Lemma 4.10. *Following the design of gradient vector field for the $(d-1)$-interface, if the deletion of regular sinks of d-interface graph disconnects the d-interface, then every connected component has at least one simplex with facet degree smaller than $d+1$.*

Proof. From Lemma 4.9 we know that the d-interface is a single connected component to begin with. Suppose that the regular sinks of d-interface graph are deleted in some sequence. Suppose that γ^{d-1} is the first simplex whose deletion disconnects the d-interface graph. Then every connected component (in the traditional graph theory sense) has at least one d-simplex which is incident on γ^{d-1} and hence upon deletion of γ^{d-1}, every connected component will have at least one simplex with facet degree smaller than $(d+1)$. The same argument can be continued for subsequent deletions and resulting disconnections. □

To see that Lemma 4.9 is essential for Lemma 4.10 to work, we see an example in Fig. 7 where lack of connectivity in the d-interface leads to components (in the d-interface) with minimum facet degree equal to $(d+1)$.

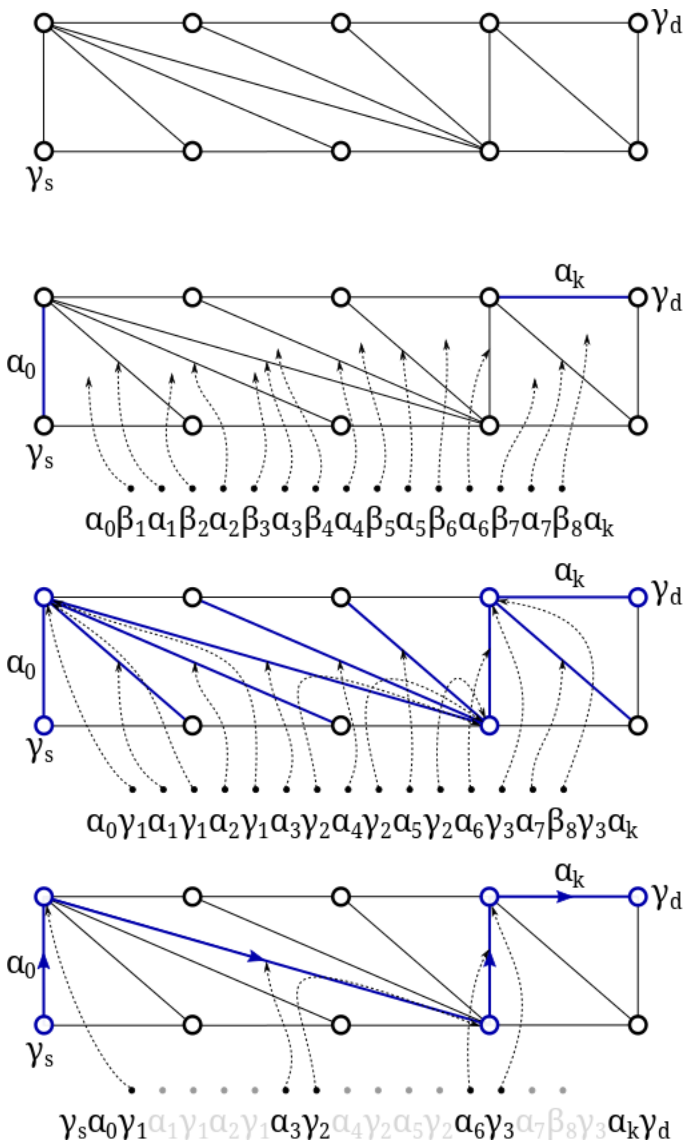

Figure 6: In this figure, we wish to establish the connectivity of γ_s and γ_d in the 1-interface. Let α_0 and α_k be 1-simplices containing γ_s and γ_d respectively. It is known that the 2-interface is connected. So, we can find the gradient sequence $\alpha_0\beta_1\alpha_1 \dots \beta_8\alpha_k$. If we let $\gamma_i = \alpha_{i-1} \cap \alpha_i$, then we can extract a new sequence $\alpha_0\gamma_1\alpha_1 \dots \gamma_8\alpha_k$. Finally, as explained in Lemma 4.8, this sequence can be used to obtain subsequence $\gamma_s\alpha_0\gamma_1\alpha_3\gamma_2\alpha_6\gamma_3\alpha_k\gamma_d$ which establishes connectivity between γ_s and γ_d.

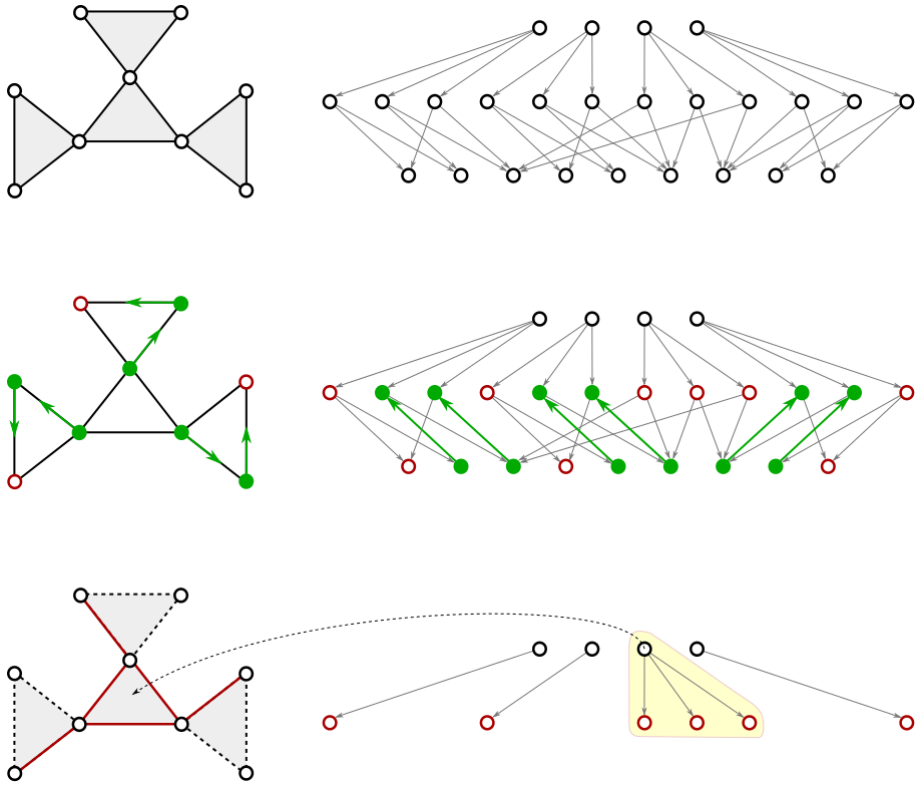

Figure 7: Unlike in the case of manifolds without boundary, in this example we have a complex whose 2-interface is disconnected to begin with. After designing vector field for the 1-interface, suppose we delete all the matched 1-simplices from the Hasse graph. Then there exists a connected component in the 2-interface for which all 2-simplices of that connected component has three facets (In this case, this connected component comprises of a single simplex with all three solid edges).

Lemma 4.11. *For a d-interface, every min-facet component has facet degree bounded by d.*

Proof. We prove the claim by induction.
Base Case: From Lemma 4.10, having deleted all the regular sinks of the d-interface, there exists at least one simplex with facet degree bounded by d in every connected component. We arbitarily choose a min-facet simplex in one of the connected components of the d-interface and discover the min-facet component around it by exploring neighboring d-simplices iteratively. Such a min-facet component has facet degree bounded by d. We design vector field on this min-facet component and subsequently delete it from the d-interface.
Induction Hypothesis: Suppose that we have designed a vector field on $(i-1)$ min-facet components and subsequently deleted them. Each time there exists at least one simplex with facet degree bounded by d in every connected component. Induction Step: Now we discover the i^{th} min-facet component say $\mathcal{F}_i$. Suppose this min-facet component belongs to some connected component $\mathcal{C}_j$.
Case 1: If $\mathcal{F}_i$ consists of all vertices in $\mathcal{C}_j$ then after processing it and deleting its vertices the other connected components continue to satisfy the facet degree (bounded by d) condition. So there is nothing to prove.
Case 2: If $\mathcal{F}_i$ consists of all the $(d-1)$-simplices of $\mathcal{C}_j$. Then upon deletion of $\mathcal{F}_i$, all d-simplices of $\mathcal{C}_j \setminus \mathcal{F}_i$ will have zero facet degree and we will attempt to match the d-simplices of $\mathcal{C}_j \setminus \mathcal{F}_i$ for the $(d+1)$-interface.
Case 3: Suppose if $\mathcal{F}_i \subsetneq \mathcal{C}_j$ and $\mathcal{C}_j \setminus \mathcal{F}_i$ has one or more $(d-1)$-simplices. Clearly there exists at least one d-simplex say σ in $\mathcal{C}_j \setminus \mathcal{F}_i$ with at least one edge incident on a $(d-1)$-simplex in $\mathcal{F}_i$ and at least one edge incident on a $(d-1)$-simplex in $\mathcal{C}_j \setminus \mathcal{F}_i$. Having designed a gradient vector field on $\mathcal{F}_i$, we delete the regular simplices and the critical $(d-1)$ simplices belonging to $\mathcal{F}_i$. Now we consider two subcases that are illustrated in Fig. 5
Case 3a: Consider the case when $\mathcal{C}_j$ stays connected after deleting the i^{th} min-facet component. In this case the facet degree of σ will reduce by at least 1 and hence the facet degree of σ is bounded by d. There may be other simplices in $\mathcal{F}_i \subsetneq \mathcal{C}_j$ whose facet degree may also reduce. All other connected components are unaffected. So, every component will have min-facet degree bounded by d.
Case 3b: Now consider the case where upon deletion of $\mathcal{F}_i$, $\mathcal{C}_j$ splits into several components. Imagine that we are not deleting the simplices of $\mathcal{F}_i$ all at once, but sequentially. Making an argument along the lines of Lemma 4.10, we conclude that irrespective of which connected component the min-facet component is chosen from, it will have facet degree bounded by d. □

Lemma 4.12. *An orientation of the min-facet component $\mathcal{F}_\mathcal{C}$ based on the matchings computed by procedure interApxMinFacet() is acylic.*

The proof of the Lemma 4.12 is identical to the proof of Lemma 4.1.

Lemma 4.13. *An orientation of the d-interface of output graph $\mathcal{H}_\mathcal{V}$ based on the matchings computed by procedure interApxMinFacet() is acylic.*

Proof. We prove this claim by induction. We use a condition namely the *vertex deletion criterion* which says that: For the d-interface, a $(d-1)$-simplex satisfies the vertex deletion condition if and only if all paths that go through that simplex end up in a sink.
Base Case: Suppose that we are processing the first min-facet component for the d-interface. From Lemma 4.12, we know that an orientation of edges of a min-facet component is acyclic. For this orientation, a path from any vertex in the component ends up in a sink. Therefore if we were to delete all the $(d-1)$-simplices in the min-facet component, we obey the vertex deletion criterion. If graph $\mathcal{H}_\mathcal{V}$ is oriented based on the matchings found in the first min-facet component, then it is acyclic.
Induction Step: Suppose that we have processed i min-facet components and suppose that we have used these min-facet components to orient the d-interface of $\mathcal{H}_\mathcal{V}$ and so far it is found to be acyclic. Also, the vertices deleted so far are those that have satisfied the vertex deletion condition. Now suppose we have extracted the $(i+1)^{th}$ min-facet component say $\mathcal{F}_{i+1}$. While the edges that lead to sinks maybe absent in min-facet component, $\mathcal{F}_{i+1}$, the corresponding d-simplices in output graph $\mathcal{H}_\mathcal{V}$ will have these edges. If we restrict our attention to undeleted edges, then from Lemma 4.12, the orientation of edges of $(i+1)^{th}$ min-facet component itself is acyclic i.e. all paths will lead strictly to critical sinks of $\mathcal{F}_{i+1}$. But if we look at the corresponding orientation in $\mathcal{H}_\mathcal{V}$, the paths emanating from a $(d-1)$ simplex of $\mathcal{F}_{i+1}$ will either end up in critical sinks of $\mathcal{F}_{i+1}$ (through undeleted edges) or in regular/critical sinks of $\mathcal{F}_j$ for $j < (i+1)$ (through deleted edges). In any case, all paths going from $(d-1)$-simplices of $\mathcal{F}_{i+1}$ go to sinks thereby satisfying

the vertex deletion criterion. Also designing gradient field on $\mathcal{F}_j$ does not introduce cycles in $\mathcal{H}_\mathcal{V}$. Morse matching on the d-interface is designed when all the min-facet components are processed and deleted. Since none of them introduce cycles, we say that output graph $\mathcal{H}_\mathcal{V}$ is acyclic. □

Lemma 4.14. *For the d-interface, the ratio $\frac{\#\text{matched simplices}}{\#\text{total simplices}} \geq \frac{2}{d+1}$.*

Proof. The proof is identical to that of Lemma 4.3 except for one important difference. In case of Algorithm 2, for a d-interface every d-simplex has $d+1$ facets. But according to Lemma 4.11, for a min-facet component the facet degree is bounded by d. Using the notion of canonical triplets for a min-facet component, for every gradient pair we get at most $(d-1)$ critical simplices. So the ratio $\frac{\#\text{matched simplices}}{\#\text{total simplices}} \geq \frac{2}{d+1}$ for every min-facet component. For a d-interface, every $(d-1)$-simplex is part of some min-facet component and is classifed as a regular simplex or as a critical simplex and subsequently deleted from the d-interface. Therefore the bound of $2/(d+1)$ carries over from min-facet components to d-interfaces. □

If we take the minimum for the ratio $2/d+1$ over all d such that $1 < d < D$, we get $2/D$. By Lemma 4.5, for a D-interface the ratio $\frac{\#\text{matched simplices}}{\#\text{total simplices}}$ is equal to $\frac{4}{D+3}$. Note that $4/(D+3) \geq 2/D$ for all $D \geq 3$. So the worst ratio of $\frac{\#\text{matched simplices}}{\#\text{total simplices}}$ over all d-interfaces where $1 \leq d \leq D$ is $\frac{2}{D}$. Since the optimal number of regular simplices $\leq$ total number of simplices, we get the following theorem.

Theorem 4.15. *For $D \geq 3$, Algorithm 3 provides a $2/D$-factor approximation for the Max Morse Matching problem.*

4.2.1 Approximation bound for nonmanifold complexes

Note that Algorithm 3 can be applied to non-manifold complexes as well if we apply the optimal algorithm for the 1-interface and procedure **intermediateApx**() for the remaining interfaces. To prove a bound for non-manifold complexes, we need to do a slightly different kind of analysis. We begin with a few definitions. Let T denote the set of all $(D-1)$-simplices of the Hasse graph. Let B denote the $(D-1)$-simplices that have been paired with $(D-2)$-simplices and let $|A| = |T| - |B|$. Let $\mathcal{R}_D$ denote the set of regular simplices found by Algorithm 3 at the D-interface. We now establish a relation between $|\mathcal{R}_D|$ and $|A|$.

Lemma 4.16. *$|\mathcal{R}_D| \geq \frac{2}{r}|A|$ where $r = D$ if the D-interface is connected and $r = (D+1)$ if the D-interface is not connected.*[4]

Proof. First we consider the case when the D-interface is not connected at the start of Algorithm 3. At each stage of Algorithm 3, the minimum facet-degree of a simplex is not more than $(D+1)$. Once again we use the idea of canonical triplets. Every critical $(D-1)$-simplex occurs in a unique canonical triplet. Also, every regular $(D-1)$-simplex occurs in at most D canonical triplets. So every regular $(D-1)$-simplex corresponds to a set of at most D critical $(D-1)$-simplices. Together they make up the entire set A. Hence we have $|\mathcal{R}_D| \geq \frac{2}{D+1}|A|$.
Now suppose the D-interface is connected at the start of the Algorithm. Then Lemma 4.10 and Lemma 4.11 apply and the minimum facet-degree of a simplex (in a min-facet component) is not more than D. In this case, a regular $(D-1)$-simplex occurs in at most $(D-1)$ canonical triplets. Accordingly, $|\mathcal{R}_D| \geq \frac{2}{D}|A|$. □

Now, let $\mathcal{R}$ denote the set of all regular simplices found by Algorithm 3 and and let $\mathcal{R}_L = \mathcal{R} - \mathcal{R}_D$.

Let $\mathcal{S}_{D-2}$ denote the set of vertices of the Hasse graph that belong to one of the d-interfaces where $1 \leq d \leq (D-2)$ and $\mathcal{S}_{D-1}$ denote the set of vertices of the Hasse graph that belong to one of the d-interfaces where $1 \leq d \leq (D-1)$. Let $\mathcal{S} = \mathcal{S}_{D-2} \cup B$. Finally, let $|\mathcal{S}|$ denote the cardinality of vertex set $\mathcal{S}$ and $|\mathcal{S}_{D-1}|$ denote the cardinality of vertex set $\mathcal{S}_{D-1}$.

Lemma 4.17. *$|\mathcal{R}_L| \geq \frac{2}{D}|\mathcal{S}|$*

Proof. Let $\mathcal{G}_\mathcal{S}$ be the graph induced by set $\mathcal{S}$. Note that every simplex belonging to graph $\mathcal{G}_\mathcal{S}$ occurs in some canonical triplet. In particular this happens to be true since all $(D-1)$-simplices of $\mathcal{S}$ are matched by Algorithm 3. Using Lemma 4.14 for Algorithm 3 applied to $\mathcal{G}_\mathcal{S}$, the ratio $\frac{\#\text{matched simplices}}{\#\text{total simplices}} \geq \frac{2}{D}$. In other words, we get, $|\mathcal{R}_L| \geq \frac{2}{D}|\mathcal{S}|$. □

[4]Note that the D-interface of a general non-manifold simplicial complex may or may not be connected.

Let $\mathcal{O}$ denote the cardinality of regular nodes found by optimal Morse Matching.

Lemma 4.18. $\mathcal{O} \leq |\mathcal{S}_{D-1}| + |T|$

Proof. The maximum number of simplices of the D-level that can be matched by any algorithm is bounded by |T| i.e. the total number of simplices of the $(D-1)$-level. Also the set $\mathcal{S}_{D-1}$ consists of all simplices of the Hasse graph except those that belong to the D-level. So, the optimal algorithm can not possibly match more than $|\mathcal{S}_{D-1}| + |T|$ number of simplices of the Hasse graph. □

We now consider the expression $D|\mathcal{R}_L| + r|\mathcal{R}_D|$,

$$\begin{aligned} D|\mathcal{R}_L| + r|\mathcal{R}_D| &\geq 2|S| + 2|A| \ldots\ldots\ldots\ldots\ldots\ldots \text{using Lemma 4.16 and Lemma 4.17} \\ &\geq |S| + |B| + 2|A| \ldots\ldots\ldots\ldots \text{using the fact that } B \subseteq S \\ &\geq (|S| + |A|) + (|B| + |A|) \\ &= |\mathcal{S}_{D-1}| + |T| \ldots\ldots\ldots\ldots\ldots\ldots \text{by definition} \\ &\geq O \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots \text{using Lemma 4.18} \end{aligned}$$

If the D-interface is connected we get,
$D|\mathcal{R}| = D|\mathcal{R}_L| + D|\mathcal{R}_D| \geq \mathcal{O}$ i.e. $|\mathcal{R}| \geq \frac{1}{D}\mathcal{O}$

If the D-interface is not connected we get,
$(D+1)\mathcal{R} = (D+1)|\mathcal{R}_L| + (D+1)|\mathcal{R}_D| \geq D|\mathcal{R}_L| + (D+1)|\mathcal{R}_D| \geq \mathcal{O}$ i.e. $|\mathcal{R}| \geq \frac{1}{(D+1)}\mathcal{O}$.

Therefore, for non-manifold complexes, Algorithm 3 gives a $1/D$ approximation if the D-interface is connected and a $1/(D+1)$ approximation if the D-interface is not connected.

Likewise one can obtain $1/(D+1)$ approximation bound for Algorithm 2, irrespective of whether or not the complex is connected.

5 Experimental results

All three approximation algorithms proposed in this paper are implemented in Java. In this section, we describe results of experiments comparing the algorithms proposed here with three algorithms for Morse matching, namely reduction heuristic, coreduction heuristic a naïve approximation algorithm that provides an approximation ratio of $1/(D+1)$. A prototype implementation was used to observe the practical performance of these algorithms on more than 800 complexes. We used both synthetic random datasets and complexes generated by Hachimori [11] (also used in earlier work [4]) and Lutz [21], for experiments. Random complexes were generated according to the method described by Meshulam and Wallach [22] and a variant. In the variant, we select a random number of valid d-simplices for all $1 \leq d \leq D$ instead of selecting a random number of D-simplices and inferring all faces. We refer to the the complexes generated by this variant as Type 2 random complexes. For additional details on these complexes see Section 5.6.

It is clear that the quantity $2|\mathcal{M}|$ where $|\mathcal{M}|$ is the size of maximum cardinality matching as well as the quantity $N - \Sigma\beta_i$ which is equal to the difference between number of simplices and the sum of Betti numbers, provide conservative upper bounds on the number of regular cells in the optimal Morse matching. Let $\mathcal{R}$ be the set of regular simplices generated by a Max Morse approximation algorithm. We *estimate* the quality of the approximation using the ratio $\frac{|\mathcal{R}|}{Min(2|\mathcal{M}|, N-\Sigma\beta_i)}$. Tables 1, 2, 4 and 5 list estimated approximation ratios on selected datasets. Algorithm 3 consistently provided the best ratios, always greater than 0.93 for all 300 random complexes in our dataset. For more than 450 manifolds from the Lutz dataset, Algorithm 3 reported worst estimated approximation ratio of 0.969. Algorithm 3 provided optimal estimated approximation ratio for 56% of manifolds from the Lutz dataset. These results suggest that Algorithm 3 not only provides good theoretical bounds, but also performs well practically.

In sections that follow, we first discuss a naïve approximation algorithm followed by experiments on datasets from four different sources.

5.1 A $1/(D+1)$-factor Naïve Approximation Algorithm

Consider the following approximation algorithm: Given a simplicial complex $\mathcal{K}$ compute its Hasse graph $\mathcal{H}_\mathcal{K}$. Perform cardinality matching on graph $\mathcal{H}_\mathcal{K}$ and obtain the matching based reorientation $\overline{\mathcal{H}_\mathcal{K}}$. Include all the down-edges of $\overline{\mathcal{H}_\mathcal{K}}$ in the output graph $\mathcal{H}_\mathcal{O}$.

1. Pick an arbitrary up-edge e and include it in $\mathcal{H}_\mathcal{O}$.
2. Include the reversed orientations of all the leading up-edges of e in $\mathcal{H}_\mathcal{O}$.
3. Remove up-edge e and the leading up-edges of e from $\overline{\mathcal{H}_\mathcal{K}}$
4. Repeat steps 1-3 until all up-edges of $\overline{\mathcal{H}_\mathcal{K}}$ are exhausted.

Clearly, $\mathcal{H}_\mathcal{O}$ has no cycles because none of the up-edges in $\mathcal{H}_\mathcal{O}$ has leading up-edges. Also, for every up-edge that we select, we reverse at most D up-edges. Since cardinality matching is an upper bound on optimal value of Max Morse Matching, we get an approximation ratio of $1/(D+1)$ for this algorithm.

At the outset, the ratio $(D+1)/(D^2+D+1)$ obtained by Algorithm 1 does not seem to be a significant improvement over $1/(D+1)$. However, as we shall witness in sections that follow, the estimated approximation ratios observed for the naïve algorithm are significantly worse in practice. In fact, in order to ensure that the approximation algorithms designed for Max Morse Matching problem remains relevant for applications like homology computation, scalar field topology etc., we need to design algorithms that can be shown to have good theoretical approximation and complexity bounds combined with competitive estimated approximation ratios.

5.2 Coreduction and Reduction heuristics

The coreduction heuristic for constructing Morse matchings was introduced in [12]. In this section, we briefly describe reduction and coreduction heuristics for constructing Morse matchings on simplicial complexes for the sake of completeness. Suppose we are given a simplicial complex $\mathcal{K}$. We first describe the coreduction heuristic.

Perfom the folowing steps until complex $\mathcal{K}$ is empty:

1. if there is a simplex α with a free coface β available, include the pair $\langle \alpha, \beta \rangle$ in the set of Morse matchings and delete α and β from $\mathcal{K}$.
2. if no such simplex with a free coface is available then we select a simplex γ of lowest available dimension and make it critical. We then delete γ from $\mathcal{K}$.

We now describe the reduction heuristic.

Perfom the folowing steps until complex $\mathcal{K}$ is empty:

1. if there is a simplex β with a free face α available, include the pair $\langle \alpha, \beta \rangle$ in the set of Morse matchings and delete α and β from $\mathcal{K}$.
2. if no such simplex with a free coface is available then we select a simplex γ of highest available dimension and make it critical. We then delete γ from $\mathcal{K}$.

5.3 Experiments on the Hachimori Dataset

This dataset consists of complexes downloaded from Hachimori's collection of simplicial complexes[5]. Table 1 lists the observed approximation ratios for all the algorithms. For complexes in Table 1, maximum size of $\Sigma\beta_i$ is 2. Clearly, coreduction heuristic provided the best approximation ratios for this dataset. However, Algorithm 3 reported ratios comparable to coreduction. Algorithm 3 reports optimal Morse matching for 7 of the 20 complexes in this dataset, while coreduction gives optimal result for 10 complexes.

[5] http://infoshako.sk.tsukuba.ac.jp/∼hachi/math/library/index_eng.html

Input	N	Estimated approximation ratios					
		Naïve	Algo 1	Algo 2	Algo 3	Cored	Red
2D complexes							
projective	31	0.800	**0.933**	**0.933**	**0.933**	**0.933**	**0.933**
dunce_hat	49	0.667	0.917	**0.958**	**0.958**	**0.958**	0.917
bjorner	32	0.667	0.933	**1.000**	**1.000**	**1.000**	0.867
nonextend	39	0.632	0.895	0.947	**1.000**	**1.000**	0.895
c-ns	75	0.703	0.892	**0.946**	**0.946**	**0.946**	0.865
c-ns2	79	0.615	0.897	0.974	0.974	**1.000**	0.846
c-ns3	63	0.667	0.871	**0.968**	**0.968**	**0.968**	0.903
simon	41	0.750	0.950	**0.950**	**0.950**	**0.950**	0.850
simon2	31	0.667	0.800	**0.933**	**0.933**	**0.933**	0.867
3D complexes							
poincare	392	0.651	0.933	0.954	0.979	**0.990**	0.923
knot	6,203	0.628	0.942	0.940	0.997	**1.000**	0.927
bing	8,131	0.640	0.946	0.943	0.997	**0.999**	0.933
nc_sphere	8,474	0.616	0.941	0.945	0.989	**1.000**	0.937
rudin	215	0.617	0.935	0.944	**1.000**	**1.000**	0.925
gruenbaum	167	0.663	0.928	0.928	**1.000**	**1.000**	0.904
ziegler	119	0.695	0.983	0.915	**1.000**	**1.000**	0.864
lockeberg	216	0.636	0.944	0.972	**1.000**	**1.000**	0.897
mani-walkup-C	464	0.645	0.944	0.922	**1.000**	**1.000**	0.922
mani-walkup-D	392	0.621	0.923	0.923	**0.990**	**0.990**	0.908
5D complexes							
nonpl_sphere	2,680	0.554	0.841	0.883	0.989	**0.997**	0.954

Table 1: Observed approximation ratios for Hachimori's Simplicial Complex Library. N indicates the number of simplices in the complex. Cored refers to Coreduction Algorithm, Red refers to Reduction Algorithm. For a given input, the best estimated approximation ratios across all algorithms tested are highlighted in bold.

5.4 Experiments on the Lutz Manifold Dataset

The second dataset consists of manifolds of dimensions ranging from 3 to 11. These manifolds were downloaded from a library of manifolds created by Lutz[6].

Table 2 lists approximation ratios observed for selected complexes within this dataset. For manifolds in Table 2, maximum size of $\Sigma\beta_i$ is 14 whereas the average size of $\Sigma\beta_i$ is 5.07. Coreduction heuristic provided the best approximation ratios for this dataset. However, Algorithm 3 matched the performance of coreduction heuristic for many complexes and in some cases outperformed coreduction. Also, Algorithm 3 was consistently better than reduction heusritic.

Table 3 summarizes the results obtained using Algorithm 3 for manifolds of different dimensions. We observed optimal results for 56% of the complexes. The worst approximation ratio was observed to be 0.969. The descriptions of homology groups of these complexes are also available in the library. For all the complexes, we compared the homology computed by application of Morse matching algorithm followed by boundary operator compuatation and Smith Normal Form with the ground truth. Our algorithm computes correct homolgy for all the complexes. The running time of homology computation was of the order of milli-seconds for most of these complexes.

[6]The Manifold page: http://page.math.tu-berlin.de/∼lutz/stellar/vertex-transitive-triangulations.html

Input	N	Estimated approximation ratios					
		Naïve	Algo 1	Algo 2	Algo 3	Cored	Red
3_12_13_3	192	0.649	0.936	0.957	**1.000**	**1.000**	0.926
3_12_1_6	240	0.672	0.933	0.908	**0.992**	**0.992**	0.882
3_15_11_1	390	0.649	0.948	0.974	0.984	**1.000**	0.953
4_15_2_24	810	0.610	0.898	0.911	**1.000**	**1.000**	0.935
4_15_4_1	965	0.566	0.875	0.902	0.987	**0.996**	0.919
5_15_2_12	1,350	0.565	0.862	0.896	0.990	**0.999**	0.951
5_14_3_16	1,120	0.572	0.873	0.898	0.998	**1.000**	0.959
6_15_2_2	5,130	0.516	0.801	0.841	0.987	**0.998**	0.961
6_15_2_1	1,890	0.546	0.847	0.877	0.995	**1.000**	0.957
7_14_3_4	6,272	0.499	0.768	0.820	**1.000**	0.999	0.955
8_14_2_15	9,326	0.479	0.747	0.782	**1.000**	**1.000**	0.962
9_15_4_1	21,310	0.458	0.716	0.757	0.996	**1.000**	0.961
10_14_38_1	15,038	0.460	0.716	0.754	**1.000**	**1.000**	0.960
11_15_2_1	30,846	0.443	0.688	0.737	**1.000**	**1.000**	0.961

Table 2: Observed approximation ratios for a few selected manifolds in Lutz's manifold library. N indicates the number of simplices in the complex. Cored refers to Coreduction Algorithm, Red refers to Reduction Algorithm. For a given input, the best estimated approximation ratios across all algorithms tested are highlighted in bold.

D	No. of complexes of dimension D	Avg size	Worst ratio	Avg ratio	% optimal
3	166	265	0.969	0.992	42.17
4	76	630	0.979	0.995	39.47
5	114	1,445	0.982	0.998	75.43
6	15	3,761	0.984	0.993	26.67
7	33	5,988	0.996	0.999	87.88
8	26	9,165	0.989	0.998	69.23
9	9	14,385	0.993	0.999	66.67
10	2	9,566	1.000	1.000	100.00
11	5	23,096	1.000	1.000	100.00
All	446	2,271	0.969	0.995	56.05

Table 3: In this table, we summarize the results for all the 446 complexes in Lutz manifold dataset. The results are grouped row-wise for each dimension. Avg size refers to the average number of simplices for complexes of dimension D. Worst ratio refers to worst estimated approximation ratio for Algorithm 3 whereas Avg ratio refers to average estimated approximation ratio for Algorithm 3, for complexes of dimension D. % optimal refers to percentage of complexes for which Algorithm 3 computes the optimal Morse matching.

5.5 Experiments on random complexes

We followed the method described by Meshulam and Wallach [22] to generate random complexes. These complexes contain all possible d-simplices for the given number of vertices, for $0 \leq d < D$. However, D-simplices are randomly chosen from all possible D-simplices based on probability $p(D)$. We generated two datasets of 100 complexes each. For each set, we generated a subset of 20 complexes with fixed $p(D)$, which varies from 0.1 to 0.9. The number of vertices was chosen to be 20 and 16 for the 4 and 6 dimensional datsets, respectively. In Table 4, we report the results for a single complex selected from each subset. It should be noted that Algorithm 3 performs well even for random complexes with non-trivial homology. For Algorithm 3, the worst estimated approximation ratio over 100 randomly generated 4-dimensional complexes was observed to be 0.939. For the 100 randomly generated 6-dimensional complexes it was observed to be 0.953. We observed that Algorithm 3 outperformed reduction and coreduction heuristics for this dataset.

$p(D)$	N	$\Sigma\beta_i$	Estimated approximation ratios					
			Naïve	Algo 1	Algo 2	Algo 3	Cored	Red
Random 6D								
0.1	16,036	3,862	0.515	0.746	0.784	**1.000**	**1.000**	0.958
0.3	18,324	1,574	0.477	0.739	0.765	**0.996**	0.989	0.967
0.5	20,612	740	0.437	0.672	0.704	**0.964**	0.940	0.936
0.7	22,899	3,003	0.428	0.653	0.705	**0.991**	0.981	0.951
0.9	25,188	5,292	0.425	0.634	0.710	**1.000**	0.998	0.953
Random 4D								
0.1	7,745	2,327	0.598	0.900	0.914	**0.999**	0.996	0.979
0.3	10,846	800	0.470	0.704	0.727	**0.952**	0.915	0.927
0.5	13,947	3,877	0.465	0.692	0.732	**0.992**	0.978	0.956
0.7	17,047	6,977	0.452	0.669	0.739	**1.000**	0.996	0.954
0.9	20,148	10,078	0.463	0.671	0.752	**1.000**	**1.000**	0.957

Table 4: This table lists estimated approximation ratios for a selected set of random complexes of dimensions 6 and 4. Each row represents results obtained by various algorithms for a single randomly generated instance. N indicates the number of simplices in the complex. Cored refers to Coreduction Algorithm, Red refers to Reduction Algorithm. $p(D)$ denotes the probability with which simplices of dimension D are chosen. $\Sigma\beta_i$ denotes the sum of Betti numbers. For a given input, the best estimated approximation ratios across all algorithms tested are highlighted in bold.

5.6 Experiments on Type 2 random complexes

We also used a variant of the Meshulam and Wallach [22] method for generation of random complexes, where we choose random number of d-simplices for all d. The generation of these random complexes proceed from lowest dimension to highest, and a random simplex is added to the complex only if all its facets are part of the complex. We generated a dataset containing 100 5-dimensional complexes with following parameters: number of vertices was chosen as 40, the probability of selecting a d-simplex is given by the vector $[1, 1, 0.7, 0.9, 1, 0.9]$. With these parameters we obtain complexes with non-trivial homology, as evidenced by their Betti numbers which lie in the range $[1, 0, 0 - 1, 2945 - 3658, 51 - 106, 0 - 3]$. Table 5 lists the results for five complexes selected from this dataset. The worst estimated approximation ratio over 100 randomly generated 5-dimensional complexes for Algorithm 3 was observed to be 0.989. We again observed that Algorithm 3 consistently outperformed reduction and coreduction heuristics for all the complexes in this dataset.

N	$\Sigma\beta_i$	Estimated approximation ratios					
		Naïve	Algo 1	Algo 2	Algo 3	Cored	Red
39,046	3,366	0.540	0.814	0.841	**0.991**	0.987	0.979
39,233	3,247	0.538	0.809	0.844	**0.993**	0.982	0.977
39,199	3,253	0.535	0.808	0.835	**0.992**	0.984	0.978
39,128	3,314	0.538	0.815	0.838	**0.994**	0.986	0.979
39,526	3,172	0.538	0.809	0.837	**0.991**	0.985	0.979

Table 5: This table lists estimated approximation ratios for a selected set of 5-dimensional Type 2 random complexes. Each row represents results obtained by various algorithms for a single randomly generated instance. N indicates the number of simplices in the complex. Cored refers to Coreduction Algorithm, Red refers to Reduction Algorithm. $\Sigma\beta_i$ denotes the sum of Betti numbers. For a given input, the best estimated approximation ratios across all algorithms tested are highlighted in bold.

5.7 Discussion on experimental results

For all datasets we studied, Algorithm 3 and Coreduction Algorithm outperform all other algorithms in terms of achieving best estimated approximation ratios. For the Hachimori dataset and the Lutz dataset, the coreduction algorithm fares slightly better, whereas for random datasets, Algorithm 3 does better. In general, Algorithm 3 outperforms all other algorithms for large sized complexes or when the size of $\Sigma\beta_i$ is large.

6 Discussion on complexity

Maximum cardinality bipartite matching is the primary bottleneck for all the algorithms described in this paper. Graph matching can be performed in $O(V^{1.5})$ time for Hasse graphs of simplicial complexes using Hopcroft-Karp algorithm [14]. With appropriate choice of data structures, all other procedures of all three Algorithms can be made to run in linear time.

In particular, for Algorithm 3, we maintain separate queues for every facet-degree. Consider the graph $\mathcal{G}$ induced by the min-facet degree simplices. To extract a min-facet component, we simply find a single connected component within this graph $\mathcal{G}$. Once the min-facet component is deleted from the d-interface, we update the facet-degrees of all affected simplices within the d-interface. Extraction and maintenance of min-facet components is therefore a linear time operation.

Also, for Algorithms 2 and 3, the approximation ratios do not depend on the graph matching steps. Graph matching step merely serves the purpose of heuristic improvement. So, effectively by removing graph matching step from Algorithms 2 and 3 become linear time approximation algorithms. But this improvement in computational complexity is at the cost of estimated approximation ratios observed in practice.

7 Conclusions and further work

We believe that approximation algorithms is the definitive algorithmic way to study Morse matchings. Our belief is validated by theoretical results and additionally supported by experimental results where we get close to optimal ratios even for random complexes.

In future, we plan to further improve the approximation bounds, remove dependency on graph matching (for improving estimated approximation ratios) and develop efficient C++ implementations. In particular, to obtain dimension independent bounds for Max Morse Matching Problem remains a challenging open problem.

Acknowledgements

The first author is grateful to Michael Lesnick for help and support. The work is partially supported by the Department of Science and Technology, India under grant SR/S3/EECE/0086/2012.

Appendix A Optimal algorithms for 1-interface and D-interface

The optimal algorithm for D-interface relies on the Lemma A.1 which has previously been proved using graph theoretic methods [15, 4]. See Appendix A of [4].

Lemma A.1. *Suppose that following the design of Morse matching for the $(D-1)$-interface, all the $(D-1)$-simplices that are matched are deleted. Then upon execution of this operation the D-interface stays connected.*

To design the vector field for 1-interface one may use DFS in the following way. Pick an arbitrary vertex s as the start vertex and mark it critical. Then invoke the procedure DFS(s):

procedure **DFSoptimal**(v,$\mathcal{G}$)

1. Mark v as visited
2. If there exists an edge $\langle v, w\rangle$ such that w is not visited then
   (a) match $\langle w, \langle v, w\rangle\rangle$
3. DFS(w)

Lemma A.2. *There exist simple linear time algorithms to compute optimal Morse Matchings for 1-interface and D-interface of D-dimensional manifolds.*

Proof. Since the graph is connected, every vertex will be visited. Also, except for the start node, every other node is matched before it is visited. The edges that are matched belong to the DFS search tree and hence do not form a cycle. Therefore, the only critical vertex is the start vertex. Therefore the simple procedure **DFSoptimal**() can be used to design optimal gradient vector field for the 1-interface. Note that the direction of gradient flow for the 1-interface will be exactly opposite of the direction of DFS traversal. We can associate a dual graph to the D-interface. Also from Lemma A.1, we know that following the design of Morse matching for the $(D-1)$-interface and deletion of matched $(D-1)$ simplices, the dual graph remains connected. So, once again, we can use the procedure on the dual graph. Therefore, upon application of DFS algorithm for the D-interface, we will have exactly one critical D-simplex which is the start vertex for the DFS and all other D-simplices are regular. This algorithm is optimal since input complex $\mathcal{K}$ is a manifold without boundary and hence must have at least one critical D-simplex. □

References

[1] Allili, M., and Kaczynski, T., Landi, C., and Masoni, F. A New Matching Algorithm for Multidimensional Persistence. newblock *arXiv preprint arXiv:1511.05427, 2015.*

[2] Ayala, R., Fernández-Ternero, D., and Vilches, J. A. *Perfect discrete Morse functions on 2-complexes.* Pattern Recognition Letters 33(11) *(2012).*

[3] Bauer, U., Lange, C., and Wardetzky, M. *Optimal topological simplification of discrete functions on surfaces.* Discrete & Computational Geometry 47*, 2 (2012), 347–377.*

[4] Burton, B. A., Lewiner, T., Paixão, J., and Spreer, J. *Parameterized complexity of discrete Morse theory.* CoRR abs/1303.7037 *(2013).*

[5] Carlsson, G. *Topology and data.* Bulletin of the english Mathematical Society *(2009).*

[6] Cazals, F., Chazal, F., and Lewiner, T. *Molecular shape analysis based upon the Morse-Smale complex and the Connolly function. In* Proceedings of the nineteenth annual symposium on Computational geometry *(2003), ACM, pp. 351–360.*

[7] Chari, M. K. *On discrete Morse functions and combinatorial decompositions.* Discrete Mathematics 217*, 1-3 (2000), 101–113.*

[8] Egecioglu, O., and Gonzalez, T. F. *A computationally intractable problem on simplicial complexes.* Comput. Geom. 6 *(1996), 85–98.*

[9] Forman, R. *Morse theory for cell complexes.* Advances in Mathematics 134*, 1 (Mar. 1998), 90–145.*

[10] Forman, R. *A user's guide to discrete Morse theory.* Séminaire Lotharingien de Combinatoire B48c *(2002), 1–35.*

[11] Hachimori, M. *Simplicial complex library.* *http://infoshako.sk.tsukuba.ac.jp/~hachi/math/library/index_eng.html , 2001.*

[12] Harker, S., Mischaikow, K., Mrozek, M., Nanda, V., Wagner, H., Juda, M., and Dlotko, P. *The efficiency of a homology algorithm based on discrete Morse theory and coreductions. In* Proc. of 3rd Intl. Workshop on CTIC *(2010), vol. 1(1).*

[13] Hersh, P. *On optimizing discrete Morse functions.* Advances in Applied Mathematics 35*, 3 (Sept. 2005), 294–322.*

[14] Hopcroft, J. E., and Karp, R. M. *An n5/2 algorithm for maximum matchings in bipartite graphs.* SIAM J. Comput. 2*, 4 (1973), 225–231.*

[15] Joswig, M., and Pfetsch, M. *Computing optimal discrete Morse functions.* SIAM J. Discrete Math. 20*, 1 (2004), 11–25.*

[16] Kaczynski, T., Mischaikow, K., and Mrozek, M. Computational Homology (Applied Mathematical Sciences)*, 1 ed. Springer, Jan. 2004.*

[17] Kozlov, D. Combinatorial Algebraic Topology*, vol. 21. Springer Science & Business Media, 2007.*

[18] Lewiner, T. *Constructing discrete Morse functions. Master's thesis, Department of Mathematics, PUC-Rio, july 2002.*

[19] Lewiner, T., Lopes, H., and Tavares, G. *Optimal discrete Morse functions for 2-manifolds.* Comput. Geom. 26*, 3 (2003), 221–233.*

[20] Lewiner, T., Lopes, H., and Tavares, G. *Toward optimality in discrete Morse theory.* Experimental Mathematics 12*, 3 (2003), 271–285.*

[21] Lutz, F. *Vertex-transitive triangulations.* *http://page.math.tu-berlin.de/~lutz/stellar/vertex-transitive-triangulations.html , 2011.*

[22] Meshulam, R., and Wallach, N. *Homological connectivity of random k-dimensional complexes.* Random Struct. Algorithms 34*, 3 (May 2009), 408–417.*

[23] Miller, E., Reiner, V., and Sturmfels, B. Geometric combinatorics*, vol. 13. english Mathematical Soc., 2007.*

[24] Mischaikow, K., and Nanda, V. *Morse theory for filtrations and efficient computation of persistent homology.* Discrete Comput. Geom. 50(2) *(2013), 330–353.*

[25] Mrozek, M., and Batko, B. *Coreduction homology algorithm.* Discrete & Computational Geometry 41*, 1 (2009), 96–118.*

[26] Shivashankar, N., M, S., and Natarajan, V. *Parallel computation of 2D Morse-Smale complexes.* Visualization and Computer Graphics, IEEE Transactions on 18*, 10 (Oct 2012), 1757–1770.*